Representative Delay Measurements (RDM): Facing the Challenge of Modern Networks

Joachim Fabini, Tanja Zseby, and Michael Hirschbichler
Institute of Telecommunications
Vienna University of Technology
Gusshausstrasse 25/E389
A-1040 Vienna, Austria
{firstname.lastname}@tuwien.ac.at

ABSTRACT

Network access technologies have evolved significantly in the last years. They deploy novel mechanisms like reactive capacity allocation and time-slotted operation to optimize overall network capacity. From a single node's perspective, such optimizations decrease network determinism and measurement repeatability. Evolving application fields like machine to machine (M2M) communications or real-time gaming often have strict real-time requirements to operate correctly. Highly accurate delay measurements are necessary to monitor network compliance with application demands or to detect deviations of normal network behavior, which may be caused by network failures, misconfigurations or attacks.

This paper analyzes factors that challenge active delay measurements in modern networks. It introduces the Representative Delay Measurement tool (RDM) that addresses these factors and proposes solutions that conform to requirements of the recently published RFC7312. Delay measurement results acquired using RDM in live networks confirm that advanced measurement methods can significantly improve the quality of measurement samples by isolating systematic network behavior. The resulting high-quality samples are one prerequisite for accurate statistics that support proper operation of subsequent algorithms and applications.

Categories and Subject Descriptors

C.4 [**Performance of Systems**]: Measurement techniques

Keywords

Measurements, networks, delay, one-way delay, real-time, anomaly, repeatability, timing, HSPA, LTE

1. INTRODUCTION

Network measurements can use either active or passive measurement methodologies, or combinations of both. Active methodologies generate dedicated measurement packets causing additional network load, whereas passive measurements identify, tag and measure packets that are part of existing traffic. One main benefit of active measurements is their ability to tailor the measurement traffic to specific measurement requirements, providing the input for a variety of applications. They help to steer application layer functions, support quality auditing and network planning and can be used as reference profiles for normal network operation when looking for anomalies and deviations caused by network failures, attacks and misconfigurations. Pronounced real-time characteristics of applications can increase demands on measurement accuracy. Functionality and stability of vital systems like the smart grid and other critical infrastructures will in the future depend on timely availability of real-time sensor and actuator information.

In nowadays networks demands for active measurements have increased substantially. Reasons include but are not limited to (a) faster networks, (b) asymmetric links, (c) aggregation of heterogeneous network technologies having distinct timing, scheduling and resource allocation strategies within one path, and (d) optimization functions like demand-driven resource allocation and compression algorithms.

This paper presents challenges, solutions and a practical realization to improve repeatability and representativity of active delay measurements. Main aim of the presented solutions is to improve the quality of raw data produced by delay measurements such that it captures representative network behavior over time. The findings have been used as basis for the prototypical implementation of a measurement framework named Representative Delay Measurement Tool (RDM) that enables accurate and representative one-way delay (OWD)measurements. The RDM prototype implementation is tested for various access network technologies in order to demonstrate achievable improvements for OWD measurements. However, the presented concepts are generally applicable and the presented tool implementation is just one possible realization.

1.1 Related Work and Main Contributions

Guidelines for accurate delay measurements have been published by the Internet Engineering Task Force's (IETF) IP Performance Metric (IPPM) Working Group in [1], [2], [3]. However, these RFCs target primarily deterministic, wired networks and must be extended to cope with the reactive nature of today's reactive mobile cellular networks. A very re-

cent update to RFC2330, RFC7312 [4], defines an advanced stream and sampling framework for IPPM with focus on uncertainty factors in active measurements.

Several researchers address requirements and new methods for OWD measurements. De Vito et al. [5] discuss in detail requirements for OWD measurements but miss the importance of start-time randomness. Various authors have published results of active and passive delay measurements in mobile cellular networks, including [6], [7], and [8]. Earlier work ([9], [10]) has presented methodological drawbacks of measurements in mobile cellular networks and compared HSPA delay figures for several HSPA network vendors and operators. Recent publications propose delay measurement techniques in reactive networks [11] and present a detailed discussion on theoretical and practical consequences of the randomness cancellation effect in time-slotted networks [12].

However, common to published tools, even to recent ones like [13], is their focus on statistical evaluation of measurement samples rather than on improvement of the sample acquisition process. We argue that advanced sample acquisition methodologies are essential to unleash systematic factors in communications and measurements. RDM fills this gap, striving for measurement samples that can reveal systematic, technology-, time-, and configuration-specific state changes in a network path. Time-critical applications can optimize their communication by adapting to these network conditions.

1.2 Structure of this Paper

The remainder of this paper is structured as follows: Section 2 defines challenges and requirements for representative OWD measurements. It identifies factors which can limit either repeatability or representativity of accurate delay measurements in wired and wireless access networks. Section 3 presents architectural and design decisions which can improve the measurement results and proposes an architecture and methodology as main requirement for accurate, representative delay measurements. Measurement results for various access technologies like HSPA and LTE are presented in Section 4, followed by conclusions in Section 5.

2. REPRESENTATIVE MEASUREMENTS

RDM targets acquisition of accurate and representative measurement samples for post-processing. The resulting set of measurement samples should (a) capture systematic variations of OWD for a measurement path operating within specifications while (b) allowing to infer on advanced technologies deployed along the path like, e.g., on-demand capacity allocation, time-slotting, or data compression devices. Knowledge of such mechanisms on a measurement path can beneficially support reliable monitoring and anomaly detection on a path. Alternatively, real-time applications can use this information to optimize data transmission.

Non-optimum measurement methodologies yield sample sets that include only subsets of the possible network behavior. Evaluating such sets raises a false sense of accuracy and reliability. One prominent example is the measurement of a time-slotted link's delay using the ping tool or, generalizing, using periodic streams. If the measurement stream's period is a multiple of the link's time slot period, the de-

lay of all samples in the set will appear as almost constant. As detailed in [12], in this case the delay of all samples depends on the global time when the measurement was started. This sample set therefore obscures the link's systematic delay variation, which amounts to a full link time slot period. Anticipating figures from the results section, $20\,ms$ minimum end-to-end delay for HSPA downlink is subject to a systematic $10\,ms$ penalty factor because of time-slotted operation. We consider the order of magnitude to be relevant for many real-time applications.

This leads to the following **main research question** as basis for the RDM concept and architecture: *Which advanced features and methodologies must a measurement tool implement to improve accuracy and representativity of OWD measurement samples?*

Central to the discussion is an appropriate definition of the term **representative**. Provided that the measurement path under observation operates within specification, i.e., is lightly loaded, we use the term *representative measurement results* as an umbrella term for several criteria that characterize the resulting set of measurement samples:

1. Path state and timing: the set of measurement samples should reliably capture – and unleash after proper post-processing – systematic and periodic variation of OWD values with time and network state of the measurement path.

2. Delay bounds: Measurement samples should allow an estimation of minimum and maximum OWD for specific packets along the observed path because of systematic state changes and timings of the measurement path.

3. Repeatability: When repeated, measurements should ideally yield comparable measurement results under identical conditions.

4. Continuity: small variations in input parameters and in measurement path state should ideally result in small variations of the measured delay.

5. Path bias: enable accurate OWD measurements for forward and for reverse link using round-trip measurements, even in the case when the forward link impairs on the quality and timing of measurement samples.

6. Direct feedback: applications should receive measurement results immediately and not rely on correlation of passive trace files.

The extent to which a specific set of measurement samples satisfies these criteria decides on how valuable the results are to applications and to which extent they might be usable to predict future delays.

It must be emphasized that the focus of this publication is strictly on acquisition methodologies for representative OWD measurement samples. Definition and selection of algorithms that post-process and analyze the samples produced by RDM is explicitly left outside of the scope. The

following subsections detail on the earlier-listed components of representative delay measurements and discuss their relevance and possible solutions.

2.1 Path State and Timing

Evolution of access networks has introduced substantial link and system state. Asymmetric link capacities for uplink and for downlink are common, both in wired and in wireless access networks. Depending on the specific technology, links are no longer stateless copper wires, but may store state and history at link layer and below. Access to links is governed by advanced scheduling mechanisms, decisions being commonly adopted by centralized schedulers. The network (IP) layer interface shields these decisions from higher layers and applications. Common examples include, e.g., time-slotted operation or demand-driven capacity allocation of a link.

Data transfer over spectrum-limited wireless links can be optimized by deploying payload or header compression. Variants include application-transparent lossless client-server compression of data over wireless links but also lossy server-only solutions. Examples for the latter can be found in cellular 3G networks. When mobile devices with limited capabilities (like small display) request web pages, some mobile operators reduce the amount of data on wireless links by silently removing optional HTTP headers from HTML documents and compressing size or quality of embedded JPEG pictures.

While all of these technological advances, scheduling and optimization strategies increase available overall network capacity, their deployment influences the OWD perceived or measured by single users and applications. Because of uplink/downlink asymmetry it is, in general, impossible to infer from round-trip delay results onto OWD. Even more difficult is to predict delay for links that exhibit on-demand capacity allocation, with direct impact onto OWD. Allocation strategies may vary, depending on user-specific parameters like recent user traffic or inter-packet delay, but also on global parameters like cell load, users in a cell, or time of day. The algorithms governing capacity allocation are unknown to users and often also unknown to operators.

The aggregated effect of these advanced mechanisms complicate representative measurements and delay predictions based on prior measurements. Methodological improvements can alleviate the impact of some of these local factors, while others – mainly the ones that depend on global parameters like users in a cell – are almost impossible to eliminate with reasonable effort.

2.2 Delay bounds

Many real-time systems depend on timely communication, benefiting from known delay bounds for their network connections. Continuous monitoring of delay bound using active or passive measurements can help to detect and handle network failures or anomalies. Safety-relevant systems might be forced to transition to a known safe state in case delay bounds change significantly.

As mentioned in Section 2.1, various systematic parameters influence on and determine minimum and maximum delay for a specific lightly loaded path – i.e., when the path operates within load specifications. Measurement samples should capture these systematic changes and – in the optimum case – enable evaluation algorithms to filter out systematic path behavior to ease detection of transient effects.

2.3 Repeatability

If repeated, delay measurements should ideally yield comparable measurement results under identical conditions. But identical conditions are impossible to guarantee in live environments. Factors like on-demand capacity allocation and global network capacity optimization can be prohibitive in achieving repeatability for specific paths and methodologies.

Still, one main goal of measurement methodologies must be repeatability in order to predict path delay – or, at least, provide realistic delay expectation values – for known conditions without the need to measure. The longer the predictions are valid, the more useful is the information to applications. Measurement methodologies can and must be improved with respect to repeatability by maintaining constant as many measurement parameters as possible. Main challenge is to define a combination of representative measurements that capture all facets of network behavior, and of measurements which closely reproduce the real traffic pattern and -load generated by applications on the network path. In addition, the amount of measurement traffic should be small when compared to the effective load generated by applications. Repeatable measurements may indicate trends which are valid beyond the measurement interval. For some paths, appropriate measurement parameters, methodologies and restrictions can help to predict future OWD from earlier measurements.

2.4 Continuity

Continuity is a property of network paths and of measurement methodologies. Applications can benefit from the knowledge that an observed network path satisfies the continuity criterion, i.e., that small changes in conditions result in small changes of the reported measurement value. Some network technologies are known to lack this property, most prominent example being time-slotted networks. A small variation in conditions – e.g., slightly delayed packet arrival at the time-slotted link – can result in the packet to miss a specific time slot and be required to wait for the next one. This incurs an end-to-end delay penalty of one full network period for the packet under observation. Which may be critical for some real-time applications.

Although measurements can not *change* the continuity property of a network path, it is still possible to *detect* that a path exhibits non-continuous behavior. This knowledge can help applications and monitoring tools to adapt their operation.

2.5 Path Bias

It is well-known that any traffic used by active measurements perturbs and potentially changes the state of network paths under observation. IETF RFC 2330 defines on this purpose the measurement property *conservative* to denote measurements which have little bias on links. But few researchers and methodologies consider the complementary effect: network paths may change the timing of measurement packets because of systematic or transient impairments. For instance, random start time packets from a specific source

will be heavily biased when leaving the first time-slotted link on their path. At the path egress, all packets will be synchronous with global time modulo network period. Therefore these packets may fail to capture the representative behavior of subsequent links on the path.

A particular case of this observation is the decomposition of round-trip- into OWDs. If the forward measurement path contains time-slotted links, it will consistently bias and change the timing of samples. Therefore, reflected packets, which are supposed to measure the reverse path, will be missing the initial (typically randomized) timing.

2.6 Direct Feedback

For operation of real-time systems, timely arrival of measurement results is of fundamental importance. Therefore we argue that immediate application-level end-to-end feedback on measurements like, e.g., through round-trip measurements should be preferred over passive measurements, which require separate channels and incur additional delays for conveying measurement information. Ideally, round-trip requests could deliver also feedback on OWDs if the path bias problem outlined in Section 2.5 is solved.

3. RDM ARCHITECTURE

Architectural and methodological solutions can address some of the challenges to representative OWD measurements in Section 2. In a top-down approach, the big picture of RDM operation is followed by a discussion of methodologies and how they can support representative measurements.

Central to RDM's operation is a strict decoupling of measurement stream definition and the measurement process. As first step "1. Configuration" in Figure 1, a command-line scenario generator tool computes representative measurement stream definitions – so-called scenario files – based on constraints passed as command-line parameters. These constraints can be value ranges and specific distribution names for packet payload, send times, server response behavior, and transfer rates. A scenario file describing one specific measurement stream is computed once, stored, and later on tested repeatedly for various setups and network parameter configurations. Alternatively, the scenario generator could translate existing tcpdump captures to scenario files. It is important to use capture files of unbiased traffic as template, e.g., packet traces captured at generating hosts. Traffic captured by passive monitoring in *intermediate* nodes might be biased by timing effects like randomness cancellation in intermediate links or systems.

RDM uses a client-server architecture as depicted in the lower frame labeled "2. Measurements" of Figure 1. Before measurements start, the RDM client reads the appropriate scenario file along with additional parameters, including at least the remote server's IP address. Supported optional parameters include the local interface to bind against, transport protocol choice (ICMP, UDP, TCP), measurement payload file or device, TTL values, and output formatting options. The RDM client then generates measurement traffic according to the scenario file description and sends it to the measurement server. The four symbolic diagrams (blue curves above SUT, red curves below) in Figure 1 emphasize that the system under test (SUT) can bias on the dis-
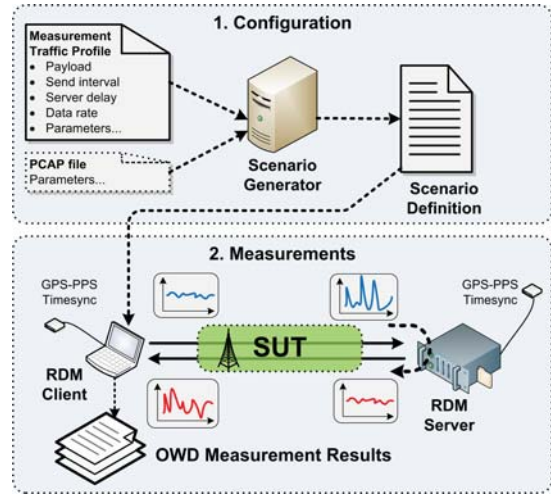

Figure 1: Overall Representative Delay Measurement (RDM) Architecture and Implementation

tribution of measurement samples while transferring them through forward and/or reverse measurement path(s), respectively. The RDM server attempts to eliminate potential measurement traffic bias of the forward path before reflecting the packets. For this it artificially delays any packet for a small, packet-specific, client-proposed value. This is why input and output distributions at the server differ substantially as illustrated by distinct diagrams in Figure 1.

When receiving a reflected measurement packet, the RDM client outputs detailed measurement results, including sequence number, absolute start time, OWD for uplink and reverse link, effective round-trip delay, requested and effective artificial server delay value, time-to-live, etc.

The remainder of this section presents the main methodological improvements that are implemented in RDM to increase representativity of measurements.

3.1 Server-based Randomness Re-Generation

RFC2330 recommends the use of random inter-packet times to avoid measurement correlation with periodic network behavior. As pointed out in Section 2.5 and in [12], it is challenging to segment round-trip delay measurement samples into OWD samples because of time-slotted *forward* link behavior. As main consequence, *reverse* link OWD samples suffer from non-random start times. Synchronization with periodic reverse link timing can originate artificial multimodal delay distributions for the reverse link.

The scatter plots in Figure 2 show this impairment effect in a live HSPA network. One point in the diagram represents the timestamp of one outbound packet modulo 100 ms function of its packet size. According to the client tcpdump report in Figure 2(a), ICMP echo requests leave the client at random send times. However, the server tcpdump arrival timestamps in Figure 2(b) identify a huge bias by a time-slotted component within the HSPA uplink which operates at 100 Hz (10 ms) clock and clusters these measurement samples.

(a) Send time (client)   (b) Arrival time (server)
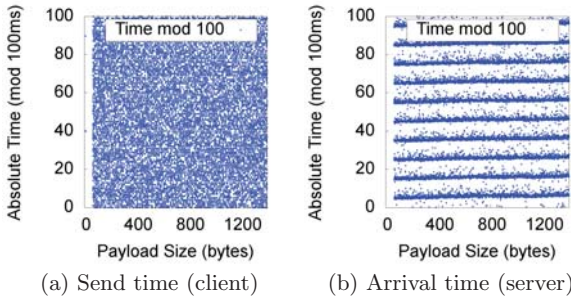
Figure 2: Effect of path bias onto random measurement samples: client send vs. server arrival timing identifies network period of 10 ms (tcpdump traces for HSPA uplink, timestamp modulo 100 ms).

Straight-forward solution to this problem is to use independent measurement streams for forward and for reverse link. RDM adopts an alternative approach which we name *server-based randomness re-generation*. For this the measurement server applies a small, artificial, random delay to any processed measurement packet before reflecting or forwarding it. The measurement server also writes incoming and outgoing timestamps to the measurement packet header, enabling the measurement client or intermediate nodes to determine accurate one-way- or hop-by-hop delay in a global-time-synchronized measurement environment. When using UDP or TCP as transport protocol, a dedicated server instance is executed on the reflecting host. For ICMP operation, the server is executed in kernel space but must support randomness re-generation extensions.
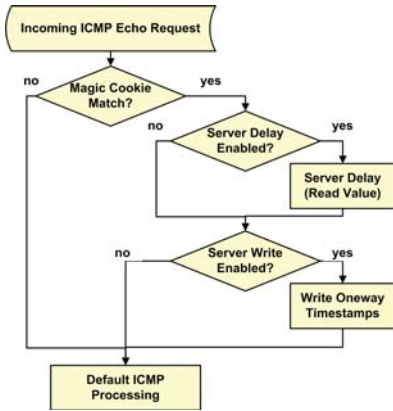
Figure 3: RDM server-based randomness re-generation and timestamping when using ICMP as protocol

The flow chart in Figure 3 illustrates the RDM server's operation. RDM client-sent measurement packets store a unique bit pattern (magic cookie) to identify their RDM protocol support. Only if the magic cookie matches, the server acts according to its internal configuration flags. If server delay functionality has been enabled, the measurement server delays any incoming measurement packet by a *client-proposed* random delay value read by the server from the measurement packet's payload field. The server writes incoming and outgoing timestamps to measurement packets when the server administrator enabled this functionality.

Main benefit of client-proposed server delay values is that the client can control measurements and configure server wait time according to its requirements and test plan for any single packet. The range of meaningful server wait times depends primarily on known or anticipated periodic network timing effects. For security reasons the RDM server wait time is limited to a maximum of 100 ms. RDM server administrators can change maximum wait limits at runtime through the Linux sysctl interface, having immediate effect on the RDM server's operation in kernel space.

The proposed randomness re-generation is extensible. Measurement packets can store several randomness re-generation headers. Cooperating hosts in the measurement path can actively support the measurement process by doing randomness re-generation for packets, each one using its own client-proposed random seed. This feature allows measuring representative OWD on a hop-by-hop basis, provided that all intermediate hosts are also time-synchronized.

3.2 Measurement Stream Definition

Measurement results in networks exhibiting on-demand capacity allocation can be heavily biased by the specific measurement stream's pattern. Even at constant average stream transfer rate during a specific observation interval, the network scheduler might allocate distinct network capacities because of distinct time distribution and payload sizes of the measurement stream's packets. Following the scheduler's decision to re-allocate network capacity for the stream, the measured delay will also change.

Listing 1: Scenario file excerpt

```
# ------------------------------------
# Generated: 2012-06-05 14:02:21
# Configured Scenario Parameters:
#    Send interval (ms):... 100-1000
#    Send interval dist:... Uni
#    Server wait (us):..... 0-9999
#    Server wait dist:..... Uni
#    Payload size (B):..... 64-1400
#    Payload dist:......... Uni
#    Total packet count:... 20000
#    Total duration (s):... 11032
#    Separator:............ <;>
# ------------------------------------
# seqNr; start time; srv wait; payload;
1;887;4504;763;
2;1595;8818;887;
3;1718;4763;552;
4;2458;7290;185;
# ...
19999;11031278;535;1045;
20000;11032025;5012;649;
```

RDM reduces this uncertainty factor by using pre-computed stream definitions as input parameter, so-called scenario files. A scenario uniquely defines a measurement stream's traffic pattern, including size, send time, artificial server delay, and, optionally, packet payload for any single stream packet. All relevant parameters and variables are computed a-priori, typically using independent random processes for computing start time and payload size. One scenario can be repeated to test and compare various network path setups and parameter settings using identical streams. The RDM client reads simple csv-separated text files as illustrated by the example in Listing 1 and translates the input scenario into an accurate measurement traffic stream according to the scenario specification, which includes reflecting server delay.

21

3.3 Supporting Representative Measurements

Server-based randomness generation and RDM's scenario concept address the requirements for representative OWD measurements as discussed in Section 2. **Server-based randomness re-generation** addresses the challenges of *path bias* (Section 2.5) by eliminating impairments that forward link timing has on measurement packets, and of *direct feeback* (Section 2.6) by storing intermediate timestamps into the measurement packet that is reflected to the sender. In addition the random server wait functionality enables clients to determine *minimum and maximum systematic delay bounds* (Section 2.2) for time-slotted reverse paths. Because of randomness re-generation in intermediate nodes, the initial sender can compute network delay of any sub-path from intermediate timestamps. Uniform distributed wait times in intermediate nodes enable statistics to compute hop-by-hop delays and their systematic variation. After eliminating outliers, lower delay bounds can be computed as the sum of minimum delays on all subpaths, whereas upper delay bounds equal the sum of the maximum delays on all subpaths. Without randomness re-generation, measurements can assess only the delay that is specific to the respective session.

RDM's **scenario concept** supports measurements in addressing their challenges in terms of *path state and timing* (Section 2.1), *delay bounds* (Section 2.2), and *repeatability* (Section 2.3). By defining appropriate measurement scenarios, measurement clients can determine to which extent a path performs on-demand capacity allocation or periodic timing. Testing identical scenario definitions in subsequent runs for different network loads lets users determine to which extent measurements in a specific environment are *repeatable*. Moreover scenarios allow to infer on a path's *continuity* property (Section 2.4) when used in lightly loaded paths.

With respect to *path state and timing* (Section 2.1), RDM's optional packet payload definition can be used to detect and outrule server-only and client-server optimizers on the path. Highly compressed measurement packet payload can outrule optimizers by preventing them to shrink the packet and benefit from lower delay. For detecting optimizers on the path, users can compare results of two subsequent runs of the same scenario, one using compressible and the other compressed measurement packet payload.

3.4 Implementation Details

Scenario generator and RDM client have been implemented in C++ using the Boost [14] and Poco [15] libraries. The RDM server's artificial randomness generation functionality has been implemented by extending the Linux kernel ICMPv4 server implementation. Ports are available for Linux kernels from 2.6 to 3.11. RDM client and server run on Ubuntu 12.04 standard PC desktop systems and laptops using a custom-compiled Linux kernel with 1kHz kernel tick. The RDM server was connected to the public Internet using a Gigabit Ethernet interface ending in the Vienna University of Technology backbone. The client accesses a live Austrian mobile cellular network using a Huawei USB E392 HSPA/LTE modem. Global time synchronization is implemented using inexpensive EM 406A based GPS-PPS solutions for clients and for servers as proposed by [6].

4. MEASUREMENT RESULTS

Unless mentioned explicitly, measurement results in this paper rely on one low-traffic measurement scenario file consisting of 20,000 measurement packets having inter-packet send intervals uniform distributed between 100-1000 ms and payload size uniform distributed between 64-1400 bytes. Artificial server randomness re-generation value is uniform distributed between 0 and 9999 μs. Total scenario duration is 11032 seconds and average scenario data rate is 10.61 kbit/s. The RDM client reads this scenario definition and generates conforming measurement streams. The low data rate has been chosen intentionally to trigger frequent state changes in on-demand-allocating mobile network links.

Fig. 4 illustrates the impact of time-slotted randomness cancellation and on-demand capacity allocation onto the representativity of reverse link measurement results. The presented use case is equivalent to the one of round-trip measurements with a traditional ping utility and randomized start times, i.e., it matches the measurement methodology recommended by pre-RFC7312 IETF documents. The measurement setup is depicted in Figure 1, measuring **HSPA** uplink as forward link and HSPA downlink as reverse link. The 10 ms transmit time interval (TTI) used by the measured public HSPA network for uplink and for downlink is identified for uplink by the time clustering in Figure 2.

The HSPA downlink delay scatter diagram in Figure 4(a) and histogram in Figure 4(e) illustrate the negative effect of randomness cancellation on measurement sample representativity. The artificial "layering", corresponding to a multi-modal sample distribution document that decomposition of round-trip delay samples into OWD samples is *not* acceptable for representative delay measurements in time-slotted networks. After enabling the RDM artificial random server delay functionality, the downlink delay diagrams in Figure 4(b) and Figure 4(f) depict the "true" downlink delay range. These diagrams are identical to the ones obtained when reversing client and server position in the measurement setup in Figure 1 such that HSPA downlink is measured first as forward link.

The uplink delay diagram in Figure 4(c) and histogram in Figure 4(g) illustrate that delay measurement results in networks with on-demand capacity allocation depend to a large extent on the specific measurement traffic pattern. The scatter plot in Figure 4(c) shows two main horizontal "layers" which differ substantially in their delay, particularly at higher payload values. Main factor which decides on whether a measurement packet is subject to lower or higher delay is preliminary state, in particular data rate and inter-packet interval. Earlier work [9] shows that mobile operators use different scheduling- and allocation policy, such that measurement results and diagrams for distinct mobile networks differ substantially for the same scenario, i.e., identical traffic.

Fig. 4(d) shows that HSPA uplink can offer unimodal delay response, too – if fine-tuning of measurement stream characteristics is a feasible option. Almost constant inter-packet send time of 150-159ms results in the much more deterministic response, confirming the finding of previous publications ([8], [9]) that higher measurement stream rate yields more
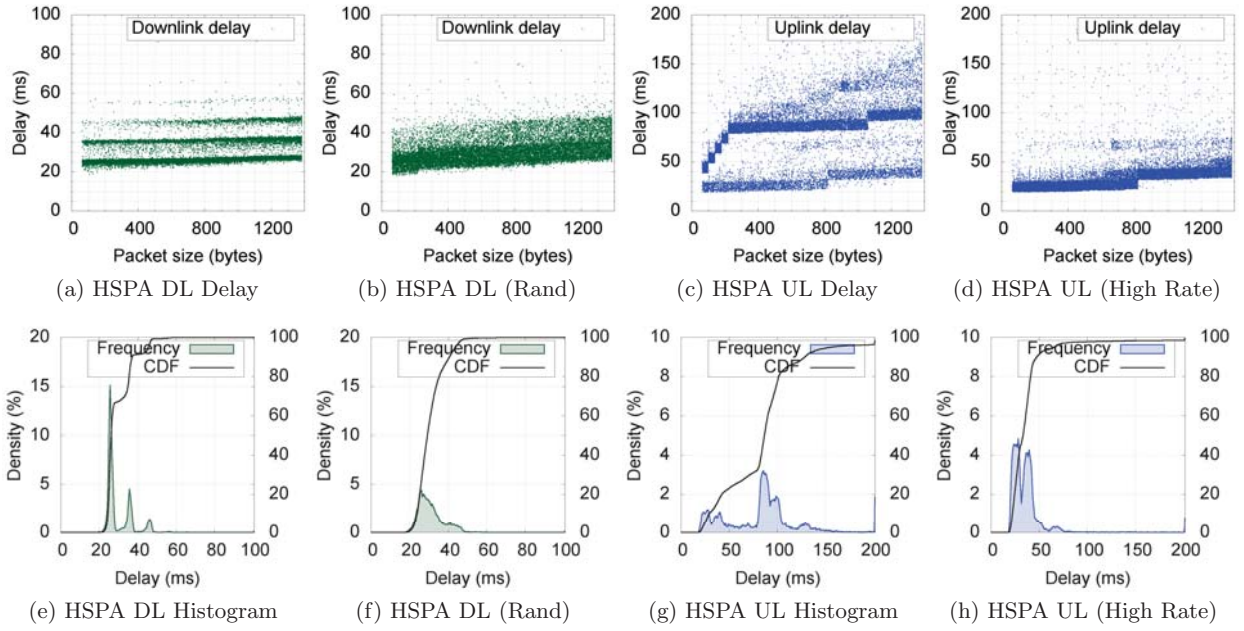
(a) HSPA DL Delay (b) HSPA DL (Rand) (c) HSPA UL Delay (d) HSPA UL (High Rate)

(e) HSPA DL Histogram (f) HSPA DL (Rand) (g) HSPA UL Histogram (h) HSPA UL (High Rate)

Figure 4: HSPA network one-way delay (DL: Downlink, UL: Uplink, Rand: Server-based randomness).

(a) LTE DL Delay (b) LTE DL Delay (Rand) (c) LTE UL Delay (d) LTE UL (High Rate)

(e) LTE DL Hist (f) LTE DL Hist (Rand) (g) LTE UL Hist (h) LTE UL (High Rate)
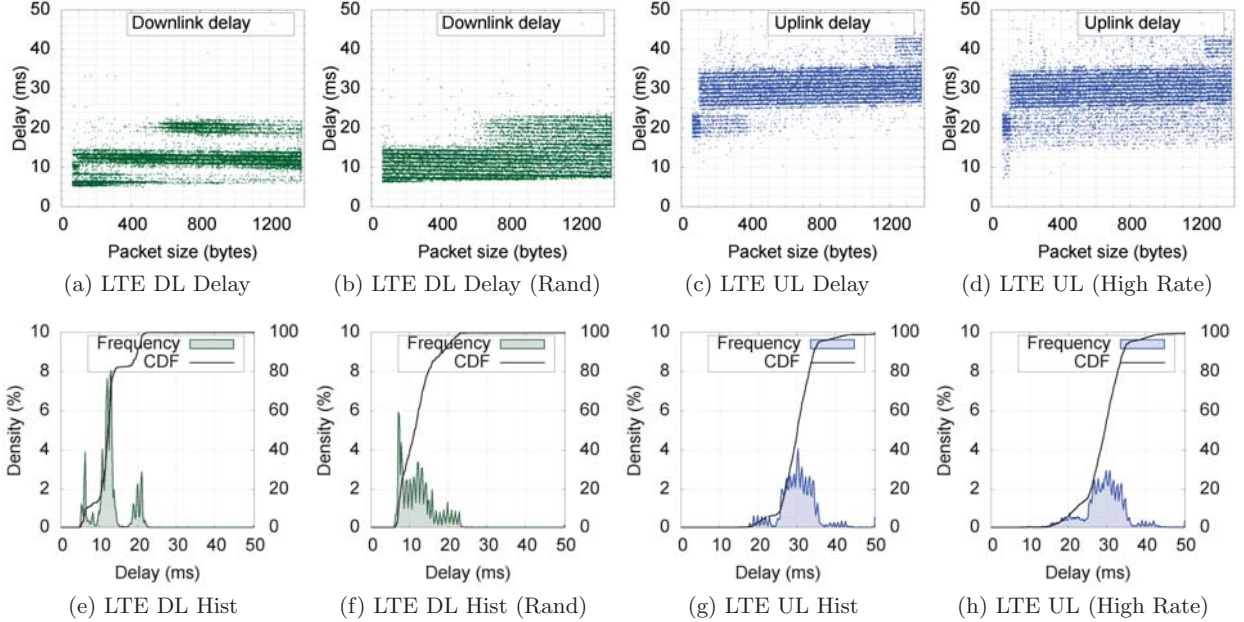
Figure 5: LTE network one-way delay (DL: Downlink, UL: Uplink, Rand: Server-based randomness)

deterministic HSPA uplink behavior.

LTE results presented in Figure 5 use the modem forced to LTE-only mode. The comparison of downlink delay diagrams in Figure 5(a) and 5(b) reveals that randomness cancellation affects LTE networks, too, and that randomness re-generation can improve the representativity of LTE results. The high-rate measurement result in Figure 5(d) shows reactive LTE uplink behavior that could be observed after a significant increase in test stream rate (10-125 ms

inter-packet interval for 64-1400 bytes payload size). Additional low-delay samples below the 25 ms limit in Figure 5(d) recommend that higher measurement rates can trigger allocation of additional capacities in the network. This knowledge can be exploited by applications. However, even if reactive behavior of the LTE scheduler can substantially improve OWD for high-rate traffic, the difference in delay – in particular of samples with higher inter-packet intervals – is by orders of magnitude lower than for HSPA.

As main limitation, when restricted to local scope, even optimum methodologies can not predict network state changes governed by global parameters and policies. The diagram in Figure 6(a) shows the result of a 35-hour measurement session (50,000 measurement packets, payload 64-1400 bytes, inter-packet send time 100ms - 5s) and demonstrates that global factors influence on HSPA uplink delay, too. The wide variety of deterministic shapes and their horizontal shift in the diagram are an indication that policies driven by cell load and time-of-day may, as well, influence on delay and contribute to the more than 400% delay variation for large payload sizes. The CDF in Figure 6(b) points out that more then 20% of all delay samples exceed the 200 ms limit. But the structured pattern of these high-delay samples, at more than five times the 90-percentile delay value of Figure 4(d), suggest that they are caused by systematic allocation effects, too.
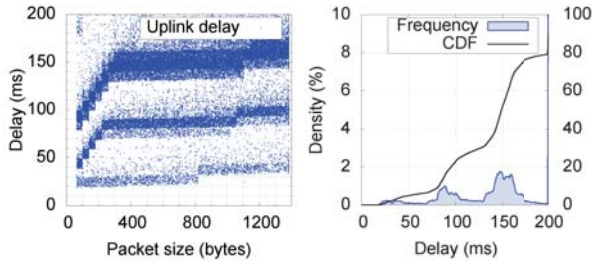
(a) HSPA UL Delay (36 hrs) (b) HSPA UL Hist (36 hours)

Figure 6: HSPA uplink delay results for 36 hours measurements.

5. CONCLUSIONS AND FUTURE WORK

Accurate and representative OWD measurements in modern access networks are highly challenging due to a variety of biasing factors, such as network state, timing effects, or reactive network capacity allocation.

This paper aims at the acquisition of representative OWD samples in modern networks. Accurate measurements support real-time applications with tight timing requirements to adapt to systematic network variations. The Representative Delay Measurement tool (RDM) is presented that implements solutions to overcome new challenges in modern access networks. Pre-computed scenario definitions support generation of identical measurement streams, whereas server-based randomness re-generation eliminates potential bias of the forward path onto measurement samples. Measurement results with RDM show that time-slotted randomness cancellation effects are observable in almost any network, the order of magnitude depending on the specific technology and configuration.

Summarizing, this paper emphasizes the importance of advanced delay sample acquisition mechanisms as prerequisite for representative statistics. We recommend that in the future all publicly available measurement data sets should be, at least, accompanied by their originating stream definitions. Servers hosting data sets should require these definition files and standardization bodies should consider to develop corresponding stream definition standards.

7. REFERENCES

[1] V. Paxson, G. Almes, J. Mahdavi, M. Mathis, *Framework for IP Performance Metrics* Network Working Group RFC 2330, May 1998.

[2] G. Almes, S. Kalidindi, M. Zekauskas, *A One-way Delay Metric for IPPM* Network Working Group RFC 2679, Sept. 1999.

[3] G. Almes, S. Kalidindi, M. Zekauskas, *A Round-trip Delay Metric for IPPM* Network Working Group RFC 2681, Sept. 1999.

[4] J. Fabini and A. Morton, *Advanced Stream and Sampling Framework for IPPM* Network Working Group RFC 7312, Aug. 2014.

[5] L. De Vito, S. Rapuano, L. Tomaciello, *One-Way Delay Measurement: State of the Art*, "Instrumentation and Measurement, IEEE Transactions on", vol.57, no.12, pp.2742 (2008).

[6] M. Laner, S. Caban, P. Svoboda, M. Rupp, *Time Synchronization Performance of Desktop Computers*, "Precision Clock Synchronization for Measurement Control and Communication (ISPCS), 2011 International IEEE Symposium on" (2011)

[7] M. Laner, P. Svoboda, P. Romirer-Maierhofer, N. Nikaein, F. Ricciato, and M. Rupp, *A Comparison Between One-way Delays in Operating HSPA and LTE Networks*, "Proceedings of the 8th International Workshop on Wireless Network Measurements WinMee'12" (2012)

[8] P. Arlos and M. Fiedler, *Influence of the packet size on the one-way delay on the down-link in 3G networks*, "Wireless Pervasive Computing (ISWPC), 2010 5th IEEE International Symposium on", pp.573-578 (2010)

[9] J. Fabini, W. Karner, L. Wallentin, T. Baumgartner, *The Illusion of Being Deterministic - Application-Level Considerations on Delay in 3G HSPA Networks*, "Networking 2009", Springer, Lecture Notes in Computer Sciences, Volume 5550, pp. 301 - 312 (2009)

[10] J. Fabini, L. Wallentin, P. Reichl, *The Importance of Being Really Random: Methodological Aspects of IP-Layer 2G and 3G Network Delay Assessment*, "Communications 2009", ICC '09. IEEE International Conference on (2009)

[11] P. Svoboda, M. Laner, J. Fabini, M. Rupp, F. Ricciato, *Packet Delay Measurements in Reactive IP Networks*, IEEE Instrumentation & Measurement Magazine, 15(6), pages 36 - 43 (2012).

[12] J. Fabini, M. Abmayer, *Delay Measurement Methodology Revisited: Time-slotted Randomness Cancellation*, Instrumentation and Measurement, IEEE Transactions on, vol.62, no.10, pp.2839-2848 (2013).

[13] R. Lübke, P. Büschel, D. Schuster, and A .Schill, *NORA: An Integrated Network Measurement Tool for End-To-End Connections*, 13th International Conference WWW/Internet (2014).

[14] BOOST C++ libraries, `http://www.boost.org/`.

[15] POCO C++ libraries, `http://pocoproject.org/`.