

Quality of service statistics over heterogeneous networks: Analysis and applications [☆]

Alessio Botta, Antonio Pescapé ^{*}, Giorgio Ventre

Università degli Studi di Napoli, "Federico II" Via Claudio 21, Napoli, Italy

Available online 24 July 2007

Abstract

Heterogeneous wireless/wired networks and ubiquitous environments are gaining ever more attention by research community. To properly control and manage such puzzles a deep knowledge of quality of service parameters is needed and, therefore, a complete and robust performance assessment is necessary. This paper deals with a performance evaluation and measurement of a number of heterogeneous end-to-end paths taking into account a wide range of statistics. To study the behavior of QoS parameters, an active measurement approach has been introduced for the analysis of properties we called (i) *concise statistics* (mean, standard deviation, inter quantile range, minimum, maximum, and median) and (ii) *detailed statistics* (Probability Density Function, Auto-correlation Function, Entropy, Complementary Cumulative Distribution Function, and Bivariate Probability Density Function). We show how, thanks to this view on QoS statistics, a more complete understanding of QoS parameters behavior is possible. In addition, we show how the measured statistics can be fruitfully used in the context of network control and management. More precisely, we present two proof of concepts regarding frameworks for QoS-based anomaly detection and for QoS-based identification of network elements.

© 2007 Elsevier B.V. All rights reserved.

Keywords: Heterogeneous networks; Test-beds and measurements; Network management systems

1. Introduction

Heterogeneous networks have gained a lot of attention in the recent past and nowadays represent a reality in almost all networking scenarios. Despite their diffusion, quality of service (QoS) provision

and analysis remain always a challenge. This is because having a well-integrated heterogeneous network infrastructure requires new approaches in several fields (e.g. planning, deploying, applications, etc.). Among the aspects to consider, performance evaluation and QoS parameter analysis play key roles during both the design and the operation phases with many interesting issues yet to be resolved.

Performance and perceived quality of *Internet applications* are primarily determined by parameters like packet loss, delay, jitter, throughput, and available bandwidth. In this direction, several reference documents containing constraints regarding these

[☆] Preliminary results within the same framework of this work have been recently published in [1].

^{*} Corresponding author. Tel.: +39 0817683856; fax: +39 0817683816.

E-mail addresses: a.botta@unina.it (A. Botta), pescape@unina.it (A. Pescapé), giorgio@unina.it (G. Ventre).

parameters have been defined [2–4], and researches have focused on the effects of these parameters on real time traffic (i.e. telephony) [5]. Moreover, in the literature there are a lot of works focused on the analysis of QoS parameters and more specifically on the delay (both one-way delay (OWD) and round trip time (RTT)) over backbone networks [6–10]. It is our opinion that, with rapidly expanding core and backbone networks, performance parameters such as delay, jitter, bandwidth, and loss get increasingly dominated by new access networks (ADSL, WiFi, etc.). Thus, the focus changes from observing performance in the network core to the communicating hosts and their access networks.

A complete understanding of the behavior of QoS parameters over heterogeneous networks is important (i) for the study of network performance, (ii) for the evaluation of network capability to support new value-added services (e.g. telephony, games), (iii) to properly design of network algorithms (routing, flow control, streaming, ...), (iv) for the definition of service level agreements (SLAs), (v) for developing algorithms to detect anomalies due to attacks, misconfigurations, measurement error, etc. and, finally, (vi) for the identification of the elements composing an end-to-end path.

To the best of our knowledge, as for the performance study, in literature a number of interesting works are present. Unfortunately, they are characterized by the use of a classical approach devoted to determine just the first-order statistics of the considered QoS parameters. At the same time, a deep knowledge of QoS statistics permits both (i) to have a complete understanding of QoS parameters behavior and (ii) to plan new and efficient control and management mechanisms. To fulfill this gap, this work has a twofold contribution.

First, we present some results regarding both the *concise statistics* (mean, standard deviation, inter quantile range, minimum, maximum, and median) and *detailed statistics* (Probability Density Function, Auto-correlation Function, Entropy, Complementary Cumulative Distribution Function, and Bivariate Probability Density Function) of the considered QoS parameters. We present such a combined statistical analysis of delay, jitter, and throughput.

Second, we show how, thanks to a deep knowledge of QoS statistics, new and efficient management strategies are possible. More precisely, we present two proof of concepts of architectures that

benefit from the use of a complete set of QoS statistics: (i) QoS-based anomaly detection and (ii) QoS-based identification of network elements.

This paper is organized as follows. In Section 2 we introduce the contributions of this work, also describing related works. Section 3 contains a description of the network test-bed and measurement tools we used. Further, it contains the description of the adopted measurement approach. In Section 4 we first present and describe the used statistical tools and then we present results of both *concise* and *detailed* statistical analysis. Also, a discussion on obtained results mainly devoted to highlight the peculiarities of heterogeneous networks is present. Finally, before ending the paper in Section 6 with concluding remarks and issues for future research, in Section 5 we present two proof of concept approaches using the QoS statistics previously analyzed in order to identify network elements and network anomalies, respectively.

2. Contribution and related works

This work has two main contributions. First, it is devoted to the characterization of QoS parameters in heterogeneous environments. Second, some applications of this study are introduced to make the reader perceive the real-world applicability of such an analysis. To highlight this twofold contribution, in this section we first present some works related to the analysis of the statistics of the QoS parameters on both wired and wireless networks, highlighting the limits of previous works. Second, we introduce some other works that use end-to-end measurement to infer both network (and host) characteristics and network anomalous behaviors due to real causes or measurement bias.

Relating to the general framework of collecting and analyzing the statistics of QoS parameters, several works have been presented in the literature. In [6] the authors measure the end-to-end delay of fixed paths demonstrating that about 84% of them present typical histograms having a Gamma-like shape with heavy tail. The authors of [8] present an analysis of the delay distribution of the SPRINT IP backbone. They find that the main contributing factor is the speed of light and that the jitter is extremely low. In [9,10] the authors measure and analyze the single-hop packet delay through operational routers in a backbone IP network. They find that it is long tailed and fits a Weibull distribution. The authors of [11] present results related to packet

delay and loss for IP data traversing the University of Auckland Internet access path. This paper contains an interesting analysis but without indication of the statistical properties of the considered variables. In [12] the authors present a jitter analysis obtained with the RIPE NCC Test Traffic Measurement setup among several measurement points located all over the world. In [13] the authors present evidence that the packet round trip delays exhibit long range dependence (LRD) showing that the complementary probability distribution decays more slowly than exponential rate. In [14] the authors study the delay characteristics of ADSL service in Korea. They measure traffic delays path by path across the whole network, to locate the bottleneck. Also, they study the relationship between delay and packet size, and between delay and network utilization. Borella et al. [15] present an analysis of the Self-similarity of Internet Packet Delay. Over some selected fixed paths, the authors find evidence that the degree of self-similarity for a round-trip path in the Internet may be correlated with the packet loss observed on that path. This result has been achieved using tools such as the variance-time plot, R/S analysis, periodogram analysis, and Whittle's estimator. As shown, there exist several qualified tools and methodologies for characterizing network behavior from end-to-end measurements. As for the statistical approach, the works present in the literature are focused on QoS parameters analysis over backbone networks and they perform a statistical analysis just in terms of CDFs and PDFs (sometimes along with the percentiles analysis). Other times they take into account the analysis of the long range dependence (LRD).

As for inferring network and host characteristics from end-to-end measurements, some interesting works have been proposed in the literature. The authors of [16] propose a passive approach aiming to detect bottlenecks located in network paths. In [17,18] two different approaches aiming to estimate the capacity of network links via end-to-end measurements are presented. In [19] the authors detect shared congestion of traffic flows using results of end-to-end measurements. The work presented in [20] is aimed at developing techniques to estimate network link loss rates by means of measures collected at a server. The work introduced in [21] presents an iterative *Bayesian* technique to identify 802.11 traffic that uses passive measurement results. Differently from [21], the work presented in [22] uses an active approach to perform the measurements.

Such work is targeted to classify the considered access networks into three categories, and it shows that it is possible to infer the belonging to a class from the results of an active probing tool. Finally, several works make use of TCP/IP properties as fingerprints able to detect some host characteristics. For instance, the authors of [23] detect host operating systems by using a Bayesian classifier.

Relating to the detection of network anomalies from active or passive measurements, the work [24] presents a framework to detect and classify, via unsupervised learning, a wide range of anomalies looking at some selected characteristics of the traffic. While, in [25], the authors introduce a framework and a set of algorithms with the aim to infer network-level anomalies from data sets of aggregate traffic.

To extend the literature, in this work we consider a more general framework in terms of both statistical approach and heterogeneous network scenarios (i.e. composed of a large mix of variables regarding the considered end-user device, operating systems, access networks technologies (Ethernet, ADSL, WLAN 802.11b, GPRS, and UMTS)). To the best of our knowledge, we extend results in the literature in that:

- we present a complete evaluation, from the application point of view, of heterogeneous end-to-end paths in terms of a wide range of QoS parameters (throughput, delay, and jitter).
- we focus our attention on a novel vision of the end-to-end path (a path including the communicating peers and their operating systems); therefore we take into account several factors like operating systems, end-user devices, network technologies and relationships among them.
- we introduce a multifarious statistical approach that includes the evaluation of different *concise statistics* (mean, standard deviation, inter quantile range, minimum, maximum, and median) as well as the study of *detailed statistics*: the probability distributions, the tail analysis, the entropy measure, the LRD analysis, and the study of the bivariate probability distribution. This permits to highlight some behavior hidden by just applying traditional statistical approach, giving insights to better understand the differences between classical and heterogeneous wired/wireless networks.
- we present some of the possible applications of our characterization in two different frameworks




Fig. 1. The experimental test-bed.

that are, the identification of the elements composing the end-to-end path and the detection of anomalies in the network due to attacks, misconfigurations, and errors in the whole chain of the measurements process.

In addition, we made publicly available the used tools and data traces at [26].

3. Test-beds, tools, and measurement approach

3.1. Real test-bed and end-to-end paths

We performed our experiments over the real test-bed described in [27] and sketched in Fig. 1. As shown, it is composed of a number of heterogeneous wireless/wired networks. Over such test-bed several configurations have been taken into account: the experiments have been performed by varying a number of configuration parameters like operating system, end-user device, access network, transport protocol, and traffic condition. All these parameters are considered as the components of our novel and wide definition of end-to-end path. In more details, in this work we define an end-to-end path (e2eP) as

$$e2eP = (S_{UD}, R_{UD}, S_{OS}, R_{OS}, S_{AN}, R_{AN}, Protocol, Bitrate) \quad (1)$$

where UD identifies the User Devices, S_{UD} at sender side and R_{UD} at receiver side (e.g. Laptop, Palmtop, Workstation, etc.); OS identifies the

Operating Systems of each of the two users (e.g. Windows, Linux, Linux Familiar¹, etc.), S_{OS} at sender side and R_{OS} at receiver side; AN is the Access Networks (LAN, 802.11, ADSL, GPRS, etc.), S_{AN} at sender side and R_{AN} at receiver side; *Protocol* is the protocol the users are communicating through (e.g. TCP, UDP, SCTP, etc.); and, finally, *Bitrate* is that imposed by the application.

By combining all these variables, in our test-bed we could set up about 350 different end-to-end paths. To highlight the features of the proposed statistical approach, as an example, we present the results of the statistical analysis performed over six end-to-end paths. The characteristics of the analyzed scenarios are presented in Table 1. TCP results have also been obtained, but in this work we focus on the UDP ones. It is worth noting that ICMP has not been used in our experiments because (i) it is handled by the routers differently from UDP/TCP and (ii) we are interested in the behavior of QoS parameters of application traffic (based on TCP and UDP).

We would like to underline that the used GPRS and UMTS connections have been provided by two of the principal Italian Telecom Operators. Such connections are the same provided to all their customers, for this reason, the reported performance are exactly the same a user would experience.

¹ An open source porting of Linux for Palmtop devices.

Table 1
Characteristics of the considered paths

ANs	Protocol	OSs	UDs
GPRS-to-Ethernet	UDP	Windows XP-to-Linux	Laptop-to-Workstation
UMTS-to-Ethernet	UDP	Windows XP-to-Linux	Laptop-to-Workstation
ADSL-to-Ethernet	UDP	Linux-to-Linux	PC-to-Workstation
Ethernet-to-GPRS	UDP	Linux-to-Windows XP	Workstation-to-Laptop
Ethernet-to-ADSL	UDP	Linux-to-Linux	Workstation-to-Desktop PC
Ethernet-to-WLAN	UDP	Linux-to-Windows XP	Workstation-to-Laptop

3.2. Measurement approach

For the active measurements we used our tool called Distributed Internet Traffic Generator (D-ITG) [28]. By combining different pairs of PS (Packet Size) and IDT (Inter Departure Time), D-ITG generates a multitude of traffic patterns. In this way we generate controlled synthetic traffic that is – at the same time – realistic. Despite this, with the aim to draw a reference curve for the parameter statistics, in this paper we consider only a UDP Constant Bitrate (CBR) traffic profile generated with constant PS and constant IDT. Besides being a traffic generator, D-ITG can be used as an active measurement tool: one-way-delay (OWD), round-trip-time (RTT), packet loss rate, jitter, and throughput, can be measured and analyzed using the various components of the D-ITG platform: (i) sender; (ii) receiver; (iii) decoder; (iv) log server. Experiments have been carried out by using three traffic conditions namely *Low*, *Medium*, and *High Traffic* [29]. For each of them, a number of PS have been used. Due to the nominal bandwidth of some of the used wireless connections (i.e. GPRS and UMTS), we consider here, only *Low Traffic* condition using IDT equal to 1/100 s and PS equal to 256 Bytes (thanks to this choice we have a maximum theoretical bit rate equal to 204.8 Kbps). To point out the end-to-end communication differences, we show the behavior of throughput, jitter, and Round Trip Time measured over UDP connections. We do not present the packet loss because, in this traffic condition, in the scenarios including WLAN and ADSL, it was always equal to 0. The jitter samples have been calculated by using the definition given in [30]. Finally, it is worth noting that the presented results have been averaged on several tests in order to minimize the effect of random error on measures. In the following, the mean values across 20 test repetitions are reported.

3.3. Collected data

The measurement stage has been performed in the time period between December 2003 and November 2004, in the day hours between 9:00 am and 6:00 pm. In that time, over 34 GB of traffic traces have been collected. Such traces have been previously inspected and sanitized in order to detect and remove samples affected by errors. At [26] we made freely available several archives containing outcomes of these measurements over real networks (which are not only those that have been used in this work). Each archive contains files with samples of QoS parameters measured over several end-to-end paths. Samples are obtained, by adopting the above described active measurement approach, sending probe packets with a rate of 100 pps and size ranging in {64, 256, 512, 1024} Bytes. More details about the traffic parameters are contained in Table 2.

Each sample is calculated using non-overlapping windows of 10 ms length. In [26] we provide also archives containing samples calculated on a per-packet basis.

3.4. Statistical methodology

By integrating and customizing established and well-known tools we have set up a methodology to provide a statistical analysis of the collected samples. Along with the evaluation of mean, standard deviation, inter quantile range, maximum, mini-

Table 2
Parameters of measurement traffic

IDT	PS	Generated bit rate
1/100 seconds	64 Bytes	51.2 Kbps
1/100 seconds	256 Bytes	204.8 Kbps
1/100 seconds	512 Bytes	409.6 Kbps
1/100 seconds	1024 Bytes	819.2 Kbps

imum values we adopted tools for the distribution estimation, study of the tails, analysis of the Auto-correlation Function, Entropy Measurement, and Bivariate Analysis. In Section 4 we present the motivations to select each of these tools in the field of heterogeneous networks. To permit the experiment repeatability and to improve the knowledge in this field, the statistical software tools (and, as said before, the data traces) that we used in this paper are freely available at [26].

4. Analysis of QoS parameters

4.1. Concise statistical analysis

Tables 3–5 present the results of the concise statistical analysis for throughput, jitter, and delay, respectively. Each table is related to one of these parameters and, for each considered path, it contains the *minimum*, the *maximum*, the *average*, and the *median* values of throughput, jitter and round trip time. Also, it contains the *standard deviation*

(StDev) and the *inter quantile range* (IQR) of the same parameters.

The IQR we used is defined as the difference between the 75th and 25th percentiles. As for these parameters, it is worth noting that average and standard deviation are more useful when analyzed along with minimum and maximum values. And, the IQR and median are better estimators for skewed distributions than, respectively, the standard deviation and the average value, because they are less influenced by extreme samples.

Digging into numerical details, as for the throughput, Table 3 shows, as expected that, in the configurations including GPRS and UMTS connections, the minimum, average, and median values of throughput are lower than those of the other configurations. Also, the standard deviation looks very similar for very different configurations. Despite this, such a result can be related to very different causes and it can be misleading if not observed together with the average values (see GPRS/ADSL-to-Ethernet). It is also interesting to note

Table 3
Concise statistics of UDP throughput [Kbps]

End-to-end path	Min	Max	Avg	StDev	IQR	Med
GPRS-to-Ethernet	0	40.96	18.44	6.913	0	20.48
UMTS-to-Ethernet	0	327.6	52.31	47.92	20.48	61.44
ADSL-to-Ethernet	122.8	245.8	204.6	6.924	0	204.8
Ethernet-to-GPRS	0	61.44	39.56	16.60	0	40.96
Ethernet-to-ADSL	20.48	348.2	204.5	15.01	0	204.8
Ethernet-to-WLAN	184.3	225.3	204.7	3.934	0	204.8

Table 4
Concise statistics of UDP jitter [s]

End-to-end path	Min	Max	Avg	StDev	IQR	Med
GPRS-to-Ethernet	0.048	5.048	0.179	0.531	0.003	0.090
UMTS-to-Ethernet	0	1.768	0.054	0.171	0.02	0.03
ADSL-to-Ethernet	0	0.089	7×10^{-4}	0.002	3×10^{-4}	7×10^{-4}
Ethernet-to-GPRS	0	0.518	0.055	0.076	0.066	0.047
Ethernet-to-ADSL	0	0.034	6×10^{-3}	0.001	4×10^{-4}	3×10^{-4}
Ethernet-to-WLAN	0	0.023	7×10^{-4}	0.001	7×10^{-4}	5×10^{-4}

Table 5
Concise statistics of UDP RTT [s]

End-to-end path	Min	Max	Avg	StDev	IQR	Med
GPRS-to-Ethernet	6.309	17.11	10.95	3.92	6.311	12.69
UMTS-to-Ethernet	1.535	4.182	2.551	0.573	0.543	2.494
ADSL-to-Ethernet	0.042	0.638	0.277	0.130	0.070	0.260
Ethernet-to-GPRS	0.801	14.31	10.15	3.017	3.040	11.26
Ethernet-to-ADSL	0.044	0.200	0.095	0.014	0.019	0.100
Ethernet-to-WLAN	3×10^{-4}	0.135	0.084	0.02	0.002	0.090

that, in the UMTS-to-Ethernet configuration, we achieved a standard deviation very close to the mean value. This implies that the average is not much representative of the sample values. That is, the samples achieved values very different from each other. In this case, we also observe a mean value quite different from the median. As for the jitter, Table 4 shows that the GPRS/UMTS based configurations achieved the worst performance (higher jitter values) also for this parameter. Indeed, they present higher maximum, average, median, and standard deviation values. In Section 4.3 a possible explanation of such behavior is present. The RTT values presented in Table 5 confirm such trend. Indeed, the values are higher for all the samples collected by using GPRS and UMTS connections.

It is worth noting that, in the case of Ethernet-to-GPRS and GPRS-to-Ethernet, the average and median values are quite different. This is not true in the case of other paths. This behavior is amplified in the case of throughput, and, it means that the role (sender or receiver) of different access networks significantly impacts on performance.

4.2. Detailed statistical analysis

In this section, applying the methodology presented in Section 3.4, we show our results on *Probability Density Functions* (PDFs) of QoS parameters as well as some results regarding the *Auto-correlation Function* (ACF), the *Entropy* measure, the *tail analysis*, and the *Bivariate Probability Density Function*. We used the Entropy to concisely quantify the randomness of the considered parameters, and the ACF to study the temporal relationship among the samples. Also, due to the fact that the distribution's behavior in its upper tail can be crucially important, we provide the tail analysis. Finally, we present the bivariate distribution of throughput and jitter to investigate the relation between these parameters.

It is important to underline that the throughput samples have been collected evaluating the average values on fixed size time intervals (100 ms), while for RTT and jitter, each packet represents one sample. Also, to plot the PDFs of all the considered parameters, we used the bin width suggested by the Scott's rule [31].

4.2.1. Probability Density Functions:

Throughput: In Fig. 2 the PDFs of the throughput samples are depicted. Such figure shows that




Fig. 2. PDF of UDP throughput.

(i) in the GPRS-to-Ethernet case, the main part (87%) of the samples achieved the median value (20.48 Kbps) while more than 10% were equal to 0 Kbps; (ii) in the UMTS-to-Ethernet scenario the samples are spread over the interval [0, 350] Kbps; (iii) in the ADSL-to-Ethernet case the median value (204.8 Kbps) has been obtained by more than 90% of the samples; (iv) in the Ethernet-to-GPRS case the samples are multi-modally distributed over 4 values (0, 20.48, 40.96, and 61.44 Kbps); (v) in the Ethernet-to-ADSL scenario even if more than 90% of the samples attained the median value (204.8 Kbps), the remaining ones range from 20.48 to 348.2 Kbps; (vi) in the Ethernet-to-WLAN case the samples are very highly concentrated around their median value (204.8 Kbps).

Jitter: In Fig. 3 the PDFs of the jitter samples are depicted. As shown, the distributions look similar in the shape, indeed, they present the majority of the samples close to 0 even if a not negligible upper tail is noticed. To better see the main part of the distribution, in Fig. 3 we included also a zoom on lower sample values. However, the sample values of the configurations including GPRS/UMTS differ of about 1 order of magnitude to the other configurations. Indeed, in the GPRS-to-Ethernet, Ethernet-to-GPRS, and UMTS-to-Ethernet cases they are mainly distributed (95% of samples) over the interval [0, 0.15] seconds (not shown). Instead, in the other cases, the 95% of the samples present values less than 0.015 seconds. It is interesting to note that the uplink of ADSL connections present quite different sample values from the downlink (the same is true also for the GPRS). Such result is due to the asymmetry of these connections. It can be used in a classification/identification framework.




Fig. 3. PDF of UDP jitter with a zoomed view.




Fig. 4. PDF of UDP RTT with a zoomed view.

Delay (RTT): In Fig. 4 the PDFs of the RTT samples are sketched. To ease the analysis, such figure includes also a zoom on the first part of the distributions. In contrast with the jitter, here the

distributions are very different from each other. Indeed, the GPRS-to-Ethernet samples are multimodally distributed around four values (7.5, 10, 12.5, and 17 seconds), some of which are not visible

in the figure. In the UMTS-to-Ethernet case the distribution is bimodal with the modes not strictly separated. In the ADSL-to-Ethernet configuration the samples are spread all over the [0.05, 0.7] seconds interval with a concentration (the 50% of the samples) around their median value (0.26 seconds). Ethernet-to-GPRS samples are close to their median value (11.26 seconds) and a heavy lower tail is present. The Ethernet-to-ADSL configuration presents samples that are mainly distributed over the interval [0.04, 0.2] seconds, and, in the mainly populated interval ([0.06, 0.12] seconds), spikes are present at multiples of 0.01 seconds. Finally, the Ethernet-to-WLAN samples are bimodally distributed around 0.02 seconds and 0.09 seconds.

Among our aims for future research there is the fitting of the depicted distributions. In the case of parameters shown in this paper, due to their shape, we plan to adopt an Expectation Maximization (EM) algorithm applied to mixture of random (probably Weibull) variables instead of techniques like model fitting and discrepancy measure [32]. EM is indeed a proper approach to easily and optimally find the parameters defining the fitting model.

4.2.2. Other statistics

ACF: To understand the samples statistical dependence, in Fig. 5 the ACF of the UDP RTT samples as a function of sample distance (called lag) is sketched. We evaluate the correlation coefficient of Pearson (r) over all lags. This is defined as

$$r = \frac{\sum_{i \in [1, n]} (X_i - \bar{X}) \cdot (Y_i - \bar{Y})}{\sqrt{\sum_{i \in [1, n]} (X_i - \bar{X})^2 \cdot \sum_{i \in [1, n]} (Y_i - \bar{Y})^2}}, \quad (2)$$

where \bar{X} and \bar{Y} represent the mean values of the two random variable X and Y . In our case, this last var-

iable (Y) is the same of X but shifted by a number of samples equal to the lag. However, in any case, it ranges from -1 to $+1$.

As the traces are related to synthetic CBR traffic (i.e. packet size and inter packet time series are constant and therefore perfectly correlated), the more the ACF values approach 0 the more uncorrelation among packet arrival times has been introduced by the end-to-end path (of course the path does not influence the packet size). As we can see, the $ACF(1)$ value is higher than 0.9 for all the considered configurations. Also, the configurations that include GPRS and UMTS connections present more uncorrelation among the samples. Indeed, for such configurations, the autocorrelation plot decays more rapidly than in the other cases. Such behavior proves that the GPRS and UMTS connections introduce uncorrelated randomness in the packet arrival process. In the case of GPRS and UMTS at sender side, the ACF shows an oscillating behavior. This is due to some periodicities in the RTT sequences. Our preliminary analysis shows that such behavior is related to the packet loss trend [33].

Entropy: In order to better understand the variability of the collected samples we have also evaluated the entropy of each trace. Generally speaking, entropy is a measure of the uncertainty of a random variable X . It is defined as

$$H(X) = - \sum_{x \in X} P(x) \cdot \log_2 P(x), \quad (3)$$

where $P(x)$ is the probability of each sample value x .

For the sake of comparing entropy values of different configurations, when estimating such parameter, we used the same bin size for all the configurations instead of that suggested by the Scott's rule. Indeed, the Scott's rule provides a bin size that varies with the samples number and values. For this reason, in the RTT case we used a fixed bin width equal to 0.01 seconds while for the jitter the bin width is equal to 0.001 seconds.

Table 6 presents entropy values calculated for the jitter and RTT samples. Such table shows that the entropy values obtained on the GPRS-to-Ethernet, Ethernet-to-GPRS, and UMTS-to-Ethernet paths are always higher than those achieved with all the other configurations. Furthermore, the reported values prove that the randomness introduced by GPRS and UMTS connections influences both the delay and its variations (jitter). Finally, it is interesting to note that when the GPRS is used at sender




Fig. 5. ACF of UDP RTT.

Table 6
Entropy of Jitter and RTT [bit]

End-to-end path	GPRS-to- Eth	UMTS-to-Eth	ADSL-to-Eth	Eth-to-GPRS	Eth-to-ADSL	Eth-to-WLAN
Jitter	3.978	3.258	0.465	6.079	0.628	0.910
RTT	4.399	6.406	2.976	8.504	2.562	1.571

side, both RTT and jitter entropy values are much higher than the other direction of the communication.

Tail analysis: In order to analyze the tail behavior we sketch the plot of the complementary CDF (CCDF) in logarithmic scales. With such a plot it turns visible the presence of a heavy tail in the distribution. It is important to underline that a close relationship exists between the delay tail behavior and the LRD of the traffic. When inputting self-similar traffics into routers, the queue length distribution exhibits a heavy-tail behavior. In Fig. 6 the CCDF of the jitter samples is depicted. In such figure it is also reported a line representing the exponential decay (it is a straight line in logarithmic axes). Therefore, such a line allows to evaluate if a distribution presents an heavy tail behavior, that is, if it has decay rate lower than exponential. With such plot it becomes clear that the jitter presents a heavy tail behavior for all the analyzed configuration.

In Fig. 7 the CCDF of RTT samples is sketched using logarithmic scales. In contrast to the previous parameter, there is no evidence of a heavy tail behavior. Indeed, for all the considered configurations, the sample distributions decay to zero with an over-exponential rate.

Bivariate analysis: In Fig. 8 the distributions of throughput and jitter are depicted, in a single plot, as the PDF of a bivariate random variable. With

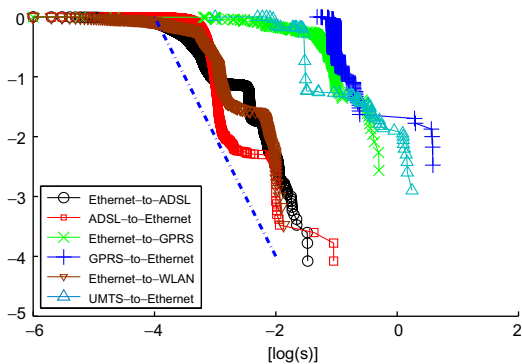


Fig. 6. Log-log CCDF of UDP jitter.




Fig. 7. Log-log CCDF of UDP RTT.

such plot it is possible to understand the relation between the distributions of the two parameters. Such figure shows that in nearly all the configurations the higher jitter values have been achieved with the median throughput. As an exception, in the Ethernet-to-GPRS case, the higher jitter values have been obtained with the lowest throughput. Further, we can observe that the highest throughput values are not associated with the highest jitter values.

4.3. Discussion

In our opinion, the approach carried out in this work underlines the importance of performing a statistical characterization based on a combined approach composed of both *concise* and *detailed* statistics. We proposed an active measurement approach using a robust platform called D-ITG. Thanks to the proposed approach interesting insights came out. As for GPRS and ADSL based configurations we have observed a relevant difference on the collected statistics depending on the communication direction. In particular, the PDFs, ACFs, and entropy values of RTT samples collected injecting traffic in the uplink direction, are different from those related to the other direction. This behavior is partially hidden when applying just a concise statistical approach. Also, a similar consideration applies for the jitter and throughput samples.

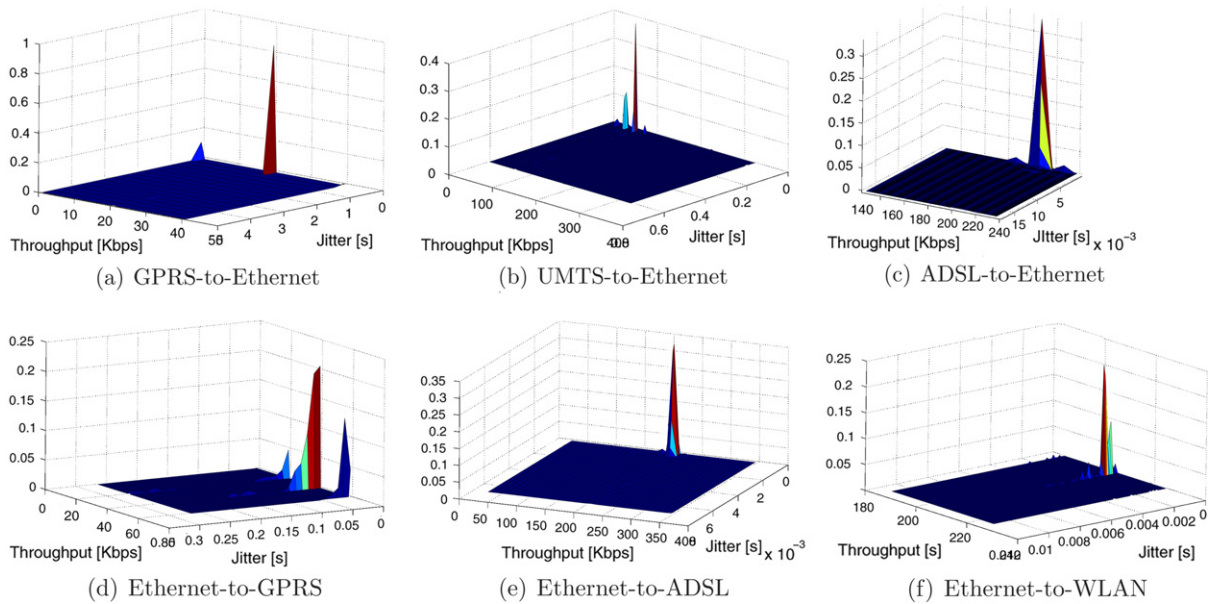


Fig. 8. Bivariate analysis of UDP throughput and jitter: (a) GPRS-to-Ethernet; (b) UMTS-to-Ethernet; (c) ADSL-to-Ethernet; (d) Ethernet-to-GPRS; (e) Ethernet-to-ADSL; and (f) Ethernet-to-WLAN.

As a general comment, we can state that there is a clear impact of the connection bandwidth on all the presented parameters and thanks to our approach we are able to completely characterize it. As an example, if we look at the RTT and jitter entropy of GPRS based configurations, we observe much higher values when the GPRS is present at receiver side. This behavior can be probably justified considering that, in this case, the packets are queued in some segment of the end-to-end path close to the receiver side. This is suitable with the low capacity of the GPRS access network. In the opposite case, with the GPRS at sender side, just the allowed packets traverse the end-to-end path and the randomness is lower than in the other communication direction.

Presented results have shown evidence that the jitter has a heavy tail behavior. Indeed, its CCDF decays with a rate lower than the exponential one. By the observation of the ACFs, we have noticed that RTT shows, in most cases, a LRD behavior. While, just for GPRS/UMTS based connections, the ACF also presents some periodicities.

As for the comparison between wireline and wireless connections, we have observed that RTT samples collected on GPRS and UMTS based configurations present the lowest values of correlation and the highest entropy values. A higher entropy value has been observed in both RTT and jitter samples. This behavior is probably due to

the fact that to transmit IP packets on a cellular network, a number of architecture elements have to be traversed before to reach the Internet. Each of these elements has its own protocol stack and gateway to the Internet. Furthermore, cellular networks have been designed to mainly transport voice traffic at 64 Kbps. Therefore, at higher bit rate, they introduce considerable latency in IP packet transmission. Also, they contribute to increase delay variation (jitter) value as already been remarked in Section 4.1.

Further analysis aiming to fully understand the driving phenomena at the base of the measured results are the subject of our ongoing work. But, on the base of current results, we are able to characterize heterogeneous wireline/wireless networks, using QoS statistics in order to propose novel approaches for network management and control.

5. Using QoS statistics to manage and control Heterogeneous Networks

In previous sections we presented a careful statistical analysis of QoS parameters giving both results and insights regarding jitter, throughput, and delay evaluated over heterogeneous wired/wireless networks. In this section we provide some example of using both *concise* and *detailed* statistics to control and manage heterogeneous networks. More precisely, we first describe a QoS-based framework

for the identification of network elements (see Section 5.1) and we then present a QoS-based framework for anomaly detection (see Section 5.2).

5.1. Identification of network elements using QoS statistics

Here we present a preliminary study we are conducting aimed to identify the component of an end-to-end path. In particular, we are working towards the development of a framework in which, by using some selected statistics of QoS parameters, it is possible to identify the single elements that compose an end-to-end communication path. As said in Section 3.1, we attribute a broad meaning to the term end-to-end path including, in its definition, also some characteristics of the end hosts. Therefore, using our approach it is possible to detect both the characteristics of the communication (*Protocol* and *Bitrate*) and the characteristics of the end hosts (*Operating System*, *User Device*, *Access Network*).

In order to decide the QoS parameter statistics to choose, a preliminary analysis has been conducted. Thanks to it we learned that the identification performs more accurately when considering the statistics contained in Table 7. In such table, r is related to the correlation coefficient of Pearson evaluated as in Eq. (2), τ refers to the correlation coefficient of Kendall, and s is the correlation coefficient of Spearman. For all of them, the number in parenthesis specify the lag at which they are calculated.

The identification framework can make use of any supervised classification algorithms. Among them we have already tried *Bayesian Networks* and *Naive Bayes*. The identification process can be summarized as follows.

First we select, for each end-to-end path, some vectors each containing values of the 13 chosen QoS parameter statistics (as in Table 7). We then partition the set of those vectors in two, obtaining a training and a test set.

Thereafter, for each component to be identified, we instruct the classification algorithm using the training set. To this aim, the vectors of this set are manually labeled with the tuple characterizing the path they were collected on. In the training phase, the classifier discovers and stores the similarities and the differences between the vectors.

Successively, we use the test set for the identification process. In this stage, the classification algorithm attributes each elements of the test set to some class. Its decision is based on what it discovered in the previous stage. Looking at the classification results, we can therefore consider a vector of statistics as correctly identified if the classification algorithm ascribes it to the correct class.

Preliminary results have proved that the identification is feasible and that high level of accuracy (i.e. the percentage of correctly identified characteristics) can be obtained ($\geq 70\%$). However, this framework is still in its early stage and several aspects have still to be investigated.

This approach can be ascribed to a “Phenomenological Identification” framework. Using QoS statistics (that depend on network behavior) we try to identify network elements.

5.2. QoS statistics-based anomaly detection systems

All the QoS parameter statistics we collected on the real test-bed of Fig. 1 allowed to deeply understand the behavior of a heterogeneous wired/wireless network. Moreover, as they were collected on a controlled test-bed, we can be confident that the characterized network was in its non-anomalous (i.e. normal) state. That is, we can state that those statistics identify the behavior of the network in its normal state. Therefore, approaches can be designed to detect an anomalous behavior of the network using those QoS parameter statistics. In more details, confronting the current values (and trends) of some selected statistics with those we collected in the normal state, we can devise if the network is behaving normally.

The network presents an anomaly when, for example, there is an attack by a group of hackers, a link is not properly working, misconfigurations occur, etc. Despite this, it may happen that the measurements indicate the network is not behaving normally even if no real causes of anomalies are present. This is the case, for example, of a problem in the measurement station or an expected but not considered route change. For this reason, in our anomaly detection research work, we are pursuing a twofold approach. The aim is to devise the best strategies to detect, on one hand, network anoma-

Table 7
Considered statistics

Mean	Median	Standard deviation	Min	Max	IQR	$r(2)$	$r(10)$	$\tau(2)$	$\tau(10)$	$s(2)$	$s(10)$	Entropy
------	--------	--------------------	-----	-----	-----	--------	---------	-----------	------------	--------	---------	---------

lies due to real causes, and on the other hand, measurement methodology errors.

In this framework, we have already started to investigate the more effective way to detect these two kinds of anomalies by looking at QoS parameter statistics. Our work is actually concerned with evaluating two different groups of factors fundamental for the detection framework, that are, the statistics to be considered and the methodologies to be utilized.

Relating to the first group, for each QoS parameter, we are investigating the capability of all the considered statistics to characterize the network behavior. This activity is carried out by injecting specific anomalies into the network and observing the variation of the different QoS parameter statistics. Clearly, depending on the type of the anomaly we inject, the most affected parameter, together with its most affected statistic, changes. Therefore, the most useful parameter (e.g. jitter, packet loss, ...) and statistic (e.g. median, IQR, ...) will be the one that guarantees, on average, the best performance for all the considered anomalies.

Furthermore, despite the optimal parameter and statistic, referring to the second group, automatic detection techniques have to be designed. In this framework, we are performing an evaluation of different methodologies able to correctly detect instantaneous value or average trend variation. The schemes we are considering are based on the following concepts: Thresholds (static or dynamic), Holt-Winters, Wavelet spectrum, and Principal Component Analysis.

6. Conclusions

All over the world heterogeneous wireline/wireless networks are being used to support mobile services and innovative multi-scenario applications. Unfortunately, even if our knowledge on their performance is increasing, a lot of work is still needed.

In this work we presented a general framework for performance characterization in real heterogeneous wireline/wireless networks from an end-user perspective. Our work extends previous works on TCP and UDP performance over WLANs in many directions. Indeed, it steps from the assumption that a current realistic scenario must consider the fusion of wired and wireless networks, several kinds of user devices, different operating systems and users' applications. More precisely, in this paper we presented some results of our empirical performance study

of a number of end-to-end heterogeneous paths, we discussed them, and finally we proposed some novel approaches to control and manage these scenarios.

A number of tests conducted on our real test-bed yielded important properties of throughput, delay, and jitter in terms of both *concise* (minimum, maximum, average, median, standard deviation, and IQR values) and *detailed* (PDF, ACF, entropy, tail analysis, and bivariate PDF) statistics. Thanks to this combined approach, several behaviors – hidden when applying just concise statistics – have been analyzed. Among them we can summarize the following: (i) we have assessed the impact of network path bandwidth on RTT and jitter and the difference between such parameters distributions for uplink and downlink traffic; (ii) we have verified that the jitter presents evidences of heavy tail; (iii) the ACF analysis of the RTT reveals LRD in all cases except the GPRS/UMTS path where we found a periodic behavior; (iv) we found highest values of entropy (of both jitter and RTT) in the case of path containing GPRS and UMTS connections.

Preliminary results have shown that the considered parameters, collected at the edge of the network, present a behavior different from the ones collected on the backbone. This suggested us to use the QoS statistics in both identification of network elements and anomaly detection frameworks.

Acknowledgements

This work has been partially supported by PRIN 2007 RECIPE Project, by Content EU NoE, and finally by Netqos and OneLab EU projects.

References

- [1] Alessio Botta, Antonio Pescapé, Giorgio Ventre, On the statistics of QoS parameters over heterogeneous networks, in: IFIP Networking 2006 Workshop Towards the QoS Internet (To-QoS'2006), May 2006, Coimbra, Portugal.
- [2] ITU-T SG16, One-way Transmission Time, Recommendation G.114.
- [3] ITU-T SG13, Network Performance Objectives for IP-based Services, Y.1541.
- [4] ITU-T SG12, The E-model, A Computational Model for Use in Transmission Planning, Recommendation G.107.
- [5] L. Zheng, L. Zhang, D. Xu, Characteristics of network delay and delay jitter and its effect on voice over IP (VoIP), ICC 1 (June) (2001) 122–126.
- [6] C.J. Bovy, H.T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal, P. Van Mieghem, Analysis of end-to-end delay measurements in Internet, in: PAM 2002.

- [7] B.-Y. Choi, S. Moon, Z.-L. Zhang, K. Papagiannaki, C. Diot, Analysis of point-to-point packet delay in an operational network, *IEEE INFOCOM* 23 (1) (2004) 1798–1808.
- [8] C. Fraleigh et al., Packet-level traffic measurements from the sprint IP backbone, *IEEE Network* 17 (6) (2003) 6–16.
- [9] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, C. Diot, Measurement and analysis of single-hop delay on an IP backbone network, *Journal on Selected Areas in Communications* 21 (6) (2003) 908–921.
- [10] K. Papagiannaki, S. Moon, C. Fraleigh, P. Thiran, F. Tobagi, C. Diot, Analysis of measured single-hop delay from an operational backbone network, *IEEE INFOCOM* 21 (1) (2002) 535–544.
- [11] K. Mochalski, J. Micheel, S. Donnelly, Packet delay and loss at the Auckland Internet access path, in: *PAM* 2002.
- [12] M. Alves et al., New measurements with the RIPE NCC test traffic measurements setup, in: *PAM* 2002.
- [13] Q. Li, D.L. Mills, On the long-range dependence of packet round-trip delays in Internet, *ICC* (1998) 1185–1191.
- [14] K.C. Kang, K. Nam, D.J. Jeon, S.Y. Kim, Delay characteristics of high-speed Internet access network: One case study, in: *APNOMS* 2003.
- [15] M.S. Borella, S. Uludag, G.B. Brewster, I. Sidhu, Self-similarity of Internet packet delay, *ICC* 1 (97) (1997) 513–517.
- [16] D. Katabi, I. Bazzi, X. Yang, A passive approach for detecting shared bottlenecks, *IEEE International Conference on Computer Communications and Networks* (2001) 174–181.
- [17] S. Katti, D. Katabi, C. Blake, E. Kohler, J. Strauss, Multiq: Automated detection of multiple bottleneck capacities along a path, *ACM Sigcomm Internet Measurement Conference* (2004) 245–250.
- [18] A.B. Downey, Using pathchar to estimate Internet link characteristics, *Measurement and Modeling of Computer Systems* (1999) 222–223.
- [19] D. Rubenstein, J. Kurose, D. Towsley, Detecting shared congestion of flows via end-to-end measurement, in: *ACM SIGMETRICS*, June 2000, pp. 145–155.
- [20] V. Padmanabhan, L. Qiu, H. Wang, Server-based inference of internet link lossiness, in: *IEEE INFOCOM* 2003, vol. 1, pp. 145–155.
- [21] W. Wei, S. Jaiswal, J. Kurose, D. Towsley, Identifying 802.11 traffic from passive measurements using iterative Bayesian inference, in: *IEEE INFOCOM* 2006.
- [22] W. Wei, B. Wang, C. Zhang, J. Kurose, D. Towsley, Classification of access network types: Ethernet, Wireless LAN, ADSL, Cable Modem or Dialup? in: *IEEE INFOCOM* 2005, Miami, pp. 1060–1071.
- [23] R. Beverly, A robust classifier for passive TCP/IP fingerprinting, in: *PAM* 2004, Juan-les-Pins, France, April 2004, pp. 158–167.
- [24] A. Lakhina, M. Crovella, C. Diot, Mining anomalies using traffic feature distributions, in: *ACM SIGCOMM* 2005, Philadelphia.
- [25] Y. Zhang, Z. Ge, A. Greenberg, M. Roughan, Network anomography, in: *IMC* 2005.
- [26] <http://www.grid.unina.it/Traffic/>. [Online].
- [27] A. Botta, D. Emma, A. Pescapé, G. Ventre, Systematic performance modeling and characterization of heterogeneous IP networks, in: *IEEE ICPADS 2005 PMW2MNC Workshop*, vol. 2, July 05, Fukuoka (Japan), pp. 120–124.
- [28] <http://www.grid.unina.it/software/ITG> [Online].
- [29] G. Iannello, A. Pescapé, G. Ventre, L. Vollero, Measuring Quality of Service parameters over heterogeneous IP networks, in: *International Conference on Networking (ICN 2005)*, LNCS, vol. 3421, 2005, April 2005, Reunion Island, 2005, pp. 718–727.
- [30] C. Demichelis, P. Chimento, IP Packet delay variation metric for IP performance metrics (IPPM), *RFC* 3393, November 2002.
- [31] D.W. Scott, On optimal and data-based histograms, *Biometrika* (1966) 605–610.
- [32] S. Pederson, M. Johnson, Estimating model discrepancy, *Technometrics* 32 (3) (1990) 304–314.
- [33] G. Iannello, F. Palmieri, A. Pescapé, P. Salvo Rossi, End-to-end packet-channel bayesian model applied to heterogeneous wireless networks, in: *IEEE Globecom 2005 General Conference*, St. Louis (MO, USA), December 2005, pp. 484–489.