Journal of Universal Computer Science, vol. 14, no. 5 (2008), 625-641
submitted: 1/11/07, accepted: 28/2/08, appeared: 1/3/08 © J.UCS

Passive Estimation of Quality of Experience

Denis Collange, Jean-Laurent Costeux
(Orange Labs, Sophia Antipolis, France
{denis.collange, jeanlaurent.costeux}@orange-ftgroup.com)

Abstract: Quality of Experience (QoE) is a promising method to take into account the users' needs in designing, monitoring and managing networks. However, there is a challenge in finding a quick and simple way to estimate the QoE due to the diversity of needs, habits and customs. We propose a new empirical method to approximate it automatically from passive network measurements and we compare its pros and cons with usual techniques. We apply it, as an example, on ADSL traffic traces to estimate the QoE dependence on the loss rate for the most used applications. We analyze more precisely the correlations between packet losses and some traffic characteristics of TCP connections, the duration, the sizes and the inter-arrival. We define different thresholds on the loss rate for network management. And we propose a notion of sensitiveness to compare these correlations on different applications.

Keywords: Quality of experience, passive performance measurements, internet performance, end-to-end performance, user behaviour, traffic analysis
Categories: C.2.3 - Network Operations, C.4 – Performance of Systems

1 Introduction

More and more new applications arise in the Internet. Some of them may be used by relatively large communities. These applications may have very different traffic characteristics and performance requirements. Internet users show then a wide range of behaviours and needs depending on the applications they use, and how they use them. These users' requirements depend also on their access mode, their terminal, etc. Some users may "suffer" on a performance level satisfying others.

An Internet Service Provider (ISP) must meet these customers' needs despite these heterogeneities. So the ISP cannot rely only on the classical network level performance criteria, such as the load and the loss rate on routers' interfaces, to manage its resources. These classical criteria only ensure that the network resources are not congested. They give very few insights on the end-to-end performance, and even less on customers' satisfaction. So a new criterion has been proposed, namely the Quality of Experience or QoE [Nokia 04] [Fiedler et al. 04]. Its aim is to characterize and to measure how the users perceive the network performance. QoE is a subjective criterion, but it must be correlated with technical network parameters managed by the network operator. It may depend on many parameters, for example the availability, the security level etc. This paper focuses on the correlations of QoE with end-to-end network level performance criteria. These correlations may be modelled by network utility functions [Fiedler et al. 05]. The challenge is then to find these correlations between the subjective QoE and the technical Quality of Service (QoS). Indeed, an ISP usually designs, monitors and manages its resources in order to respect some network level performance criteria. To be sure that its customers experience a good

QoE level, this ISP has to know the correlations between QoS and QoE. There are many ways to get these correlations.

A first one is to carry out intensive experiments with a panel of users, on a test-bed platform, varying the network performance in a controlled way and asking the users to grade their QoE. This is a quite precise method, and it has been applied to specify standards for few applications and QoS criteria: the telephony [ITU 03] and the web [ITU 05]. A drawback is its workload. It takes a lot of time to be carried out while new applications appear every month. Some of them disappear, but other applications "explode" in short time. Moreover, the existing applications frequently have new releases, with sometimes new traffic characteristics, higher performance requirements, different behaviour, and so on. So it would be too expensive to apply this first method on every application, on every new release.

A second possible method is to over-provision the network and to monitor its QoS waiting for users' complaints. When they appear, the ISP increases the insufficient resources and deduces some correlations between the network-level performance and the users' satisfaction. This method is also expensive, as the network resources are costly. Moreover, the ISP risks angry users moving to other ISP.

We propose here a third method. We suggest to analyze the customers' behaviour through many traffic characteristics and to study the correlations between these characteristics and the QoS. We can indeed assume that customers modify their behaviour when they get bothered by bad performance. They may interrupt active transfers, or avoid or postpone some large transfers. Some users may renew previously stopped transfers and then increase the arrival rate of transfers. They may also wait for the end of an active transfer before launching new ones, reducing then the arrival rate. Changes in the users' behaviour may thus be an objective consequence of degradation of the network performance. This method must be applied passively, i.e. on actual customers' flows and not on test transfers. It is rather cheap and fast, as compared to the previously mentioned methods. It can also be frequently applied following new applications and releases that constantly arise. Of course, this method is not perfectly accurate to find the exact correlations between the subjective performance perceived by the users and the network QoS. These correlations depend not only on the users' behaviour, but also on the protocols reactions at different levels, from the transport level (TCP) to the application level. For example, even a small loss rate causes TCP retransmissions increasing the duration and the volume of TCP connections. On the contrary, with high loss rate or variable round trip time, the TCP protocol or the application may abort the connections. Nevertheless our method can be used to define some thresholds on the network QoS criteria. For example, a first one when the traffic characteristics begin to be impacted by QoS degradation, a last one when the performance is so bad that the traffic features do not change anymore. It may also be used to characterize users' (or applications') sensitivity to the network level QoS.

These three methods are then not exclusive, but rather complementary. None of them can be applied on all the applications to define the exact correlations between QoE and QoS. The first method should rather be used on the main applications (most frequently used, most important for the users) to know precisely the users' requirements. The other two methods should be used simultaneously to manage networks more roughly, to analyze tendencies, to take into account new applications etc.

To start with this third method, we analyze in this paper the correlations between packet losses of TCP connections and some of their traffic characteristics. Our aim is, to show firstly the existence of these statistical correlations, and secondly the possibility to use them to define thresholds and characterize applications' sensitivity for monitoring a specific network at a given time. We do not intend to give general "worldwide applicable" correlations between the users' behaviour and the loss rate. These correlations depend on too many psychological, regional and temporal parameters. As already mentioned, we would rather think that our method is to be applied frequently on the network to be managed. Moreover, we want to show the existing correlation between the loss rate and various traffic characteristics (the duration, the size and the arrival rate). Actually those characteristics define only some aspects of the users' behaviour. We also limit this first analysis to the loss rate. We have presented an extension of this analysis with other performance criteria and new traffic traces on a bigger BAS in [Collange 07]. We have chosen to present the analysis on the end-to-end loss rate of TCP connections, as it is a more technical performance criterion. It is closer to the classical criteria used in network management. Moreover it depends essentially on the network state while, for example, latency may be impacted by application timers, throughput by rate limitations etc.

Some authors analyzed various models concerning the alteration of users' behaviour, but they do not refer to observations on operational networks [Guillemin et al. 03], [Roberts and Bonald 03] and [Yang and Veciana 01]. Former traffic analysis on the impact of the performance on users' behaviour considered the interruptions of connections [Peuhkuri 02] [Rossi 02] or the round-trip time [Collange et al. 05].

First, we present the architecture of the network, our traces, and two simple methods to approximate the end-to-end packet loss rate in Section 2. In Section 3 we show the existence of correlations between packet losses and some traffic characteristics: the means of flow duration, up and down sizes, and the arrival rate. Based on these correlations, we propose various thresholds of the loss rate for network management. We also correlate the losses with the terminations of connections. Next, we try to quantify these correlations in Section 4. We propose a method to compare the sensitiveness of applications to the performance degradation. We apply this method to the most frequently used applications to compare how their traffic features depend on the loss rate.

2 Measurements settings

In subsection 2.1 we first describe our collection infrastructure, based on a running network. Then, in subsection 2.2, we explain how we passively detect the end-to-end packet losses on our traffic traces.

2.1 Traffic capture

A classical ADSL architecture is organized as follows (see Figure 1): the BAS (Broadband Access Server) collects the traffic issued from many DSLAMs (Digital

Subscriber Line Access Multiplexers) before forwarding it through the local routers to the France Telecom IP backbone. Each client is connected to only one DSLAM.

Our probe is located between a BAS and the first routers to the Internet. This BAS multiplexes the traffic of 10 DSLAMs connecting around 4000 clients. We capture the whole TCP/IP headers of all the packets going through the BAS without any sampling or loss (using an optimized adaptation of tcpdump [Jacobson et al. 89]). The collected data thus represents a huge amount of traffic and can be a good representation of the variety of the applications currently used in the Internet. Of course, as the relative usage of the applications depends on the region, our final results are more "local" and specific. The analysis that we present here refers to one complete week of measurements in 2006.

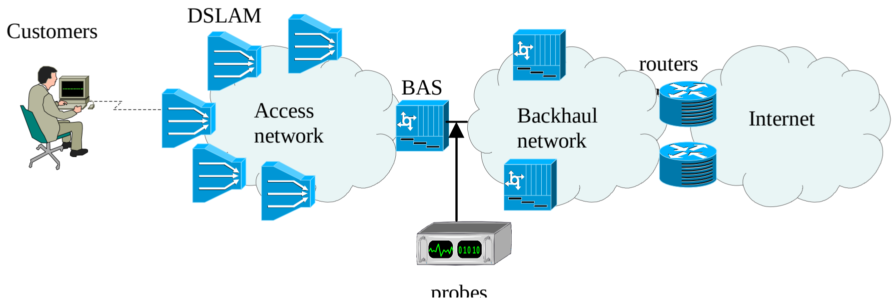

Figure 1: The architecture of the ADSL access network

2.2 Loss rate measurement

From these traffic traces we compute and store in real-time many criteria characterizing the traffic and the performance of TCP and UDP connections. We define and implement two methods to estimate the end-to-end loss rate for each TCP connection basing on the TCP/IP packet headers. This measurement is only an approximation, as we use passive measurements in the middle of the network.

The first method measures the proportion of what we call "desequencements". Packets are "desequenced" if their sequence number is lower than the sequence number of the very last packet seen in the same TCP connection. This is for example the case of the first two packets circled on Figure 2. This figure presents the history of the sequence numbers of a TCP connection where many packets were lost and retransmitted. The last three packets in circles are not desequencements. We assume that we observe a desequencement due to the retransmission (by TCP) of a lost packet, if at least one packet has been lost in the last congestion window. So, the desequencement rate usually underestimates the actual loss rate. However, some routers' behaviours may also generate desequencements without loss.

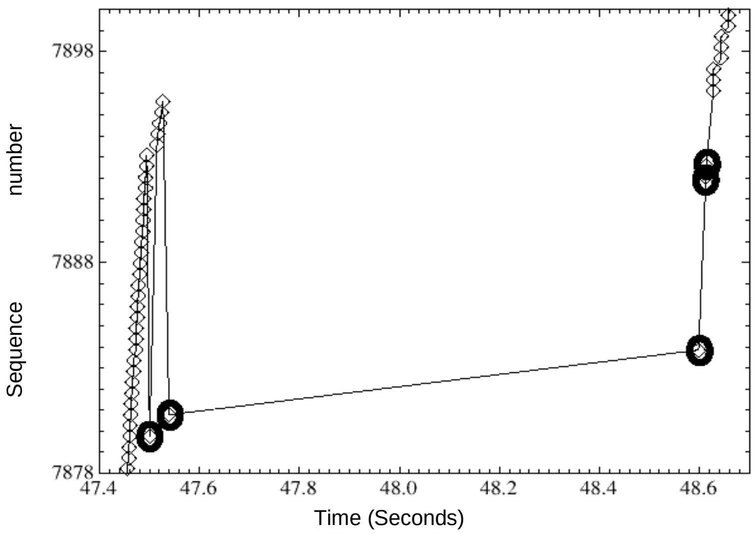

Figure 2: Desequencements and retransmissions on the history of TCP sequence numbers of a TCP connection

Another method to estimate the loss rate is based on retransmitted packets, i.e. when we see a given sequence number twice. This is the case for all circled packets on Figure 2. We call such events "retransmissions", even if some of them may be duplicated by network nodes, or if packets are lost before the probe and retransmissions are not counted. As the Selective Acknowledgments option is not always implemented, TCP may have the Go-Back-N behaviour and it may retransmit more packets than needed. Consequently, this retransmission rate may overestimate the actual loss rate, but it depends more precisely on the location of the probe.

In this paper we consider the ratio of the lost packets to the total number of packets, including retransmissions, initial synchronisation packets etc. These methods apply only to the TCP connections, UDP does not use sequence numbers. Despite this fact, they allow to detect the state of the network, since the TCP protocol represents more than 90% of the Internet traffic volume. In [Costeux and Guyard 06] we apply a similar analysis on the RTP transfers, i.e. on the UDP applications whose loss rate can be measured. This is typically the case of the voice or the video over IP.

Some more accurate methods have been described, for example in [Jaiswal et al. 03] and [Brosh et al. 05], to detect the end-to-end packet losses more precisely. In [Brosh et al. 05] the authors consider not only TCP but also IP sequence numbers. This helps to distinguish retransmissions from packets duplicated by network nodes, or losses before and after the measurement point. The other algorithm proposed in [Jaiswal et al. 03] infers the size of the congestion window and the current round-trip time from TCP segments and acknowledgments observed by the probe. It applies then six rules to distinguish between the possible causes of out-of-sequence packets. As noted in [Brosh et al. 05], this algorithm is sensitive to delay variations. Moreover, it needs a lot of CPU time to estimate the round-trip time and the congestion window at each packet of each TCP connection. For this reason, it cannot be applied in real time to thousands of simultaneous active TCP connections. Hence, we use rough, but quickly computable approximations of the loss rate. Besides, our aim is not to estimate the end-to-end loss rate accurately, but to show the existence of correlations between the loss rate and some traffic characteristics. We use two relatively different approximations: one rather over-estimating and the other rather under-estimating the effective loss rate. As we notice similar behaviour with both estimations (Figure 8), we deduce that there is a correlation between the traffic characteristics and the end-to-end loss rate no matter the method used to detect the losses.

3 Correlations between the loss rate and traffic characteristics

As observed in [Peuhkuri 02], the users are particularly sensitive to the response time of their requests, i.e. to the time to get their transactions finished. Thus we first consider in subsection 3.2 the impact of losses on transfer durations. We observe that the response time is significantly influenced by the retransmission delays of lost packets. In subsection 3.3 we consider the users' behaviour through their flow sizes. To explain the observed changes in transfer sizes according to the loss rate, we study in subsection 3.4 if the interruptions of connections can be the actual cause of small sizes in case of bad performance. Finally we analyze in subsection 3.5 the impact of losses on the arrival rate of TCP connections.

3.1 Overall traffic characteristics over the day

Our aim in this subsection is not to present a general traffic analysis but only to exclude some possible hidden factors in our following analyses, such as simultaneous daily variations.

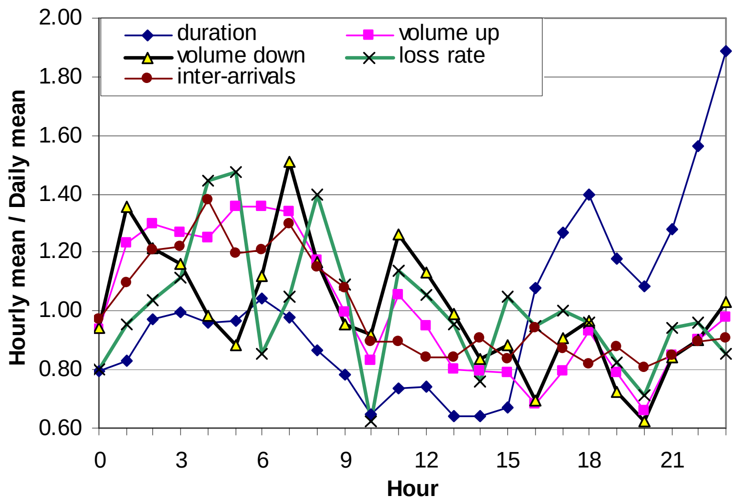

Figure 3: Daily variations of traffic characteristics

Figure 3 presents the variations of the loss rate and of several characteristics of TCP connections over one weekday. The averages per hour are divided by the daily averages in order to look at the variations on the same scale. We notice that no matter the considered feature, its mean per hour does not vary a lot (about 40% around the daily mean). We will observe in the next subsections much bigger variations, of many orders of magnitude, in the relative dependence between the traffic characteristics and the loss rate. Here the greatest variation appears with the transfer duration which doubles in the evening, without clear correlations with the other characteristics. This may indicate degradation of the transfer throughputs due to the increase of the residential load in the evening. We also notice that the transfers are larger during the night.

Figure 4 gives the proportion of connections with non-zero loss rate. It shows that the most significant part of these connections encounters loss rates greater than 1%. The high proportion of flows encountering high loss rate comes from short connections (10 to 100 packets) which lose few packets.

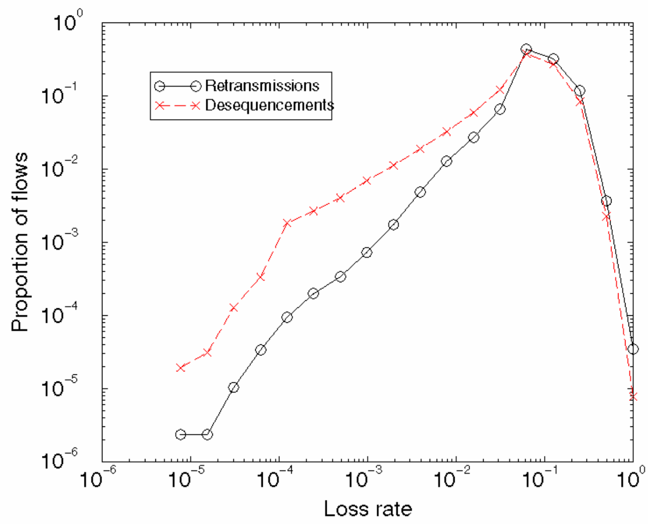

Figure 4: Proportion of flows according to the loss rate

Figure 5 shows the cumulative distribution functions of the number of connections and their upward and downward traffics, according to their size in bytes. We classically observe that 95% of the connections are very small, whereas 80% of the traffic comes from connections larger than 1 MB (Megabyte). Even 50% of the traffic is carried by connections larger than 10 MB.

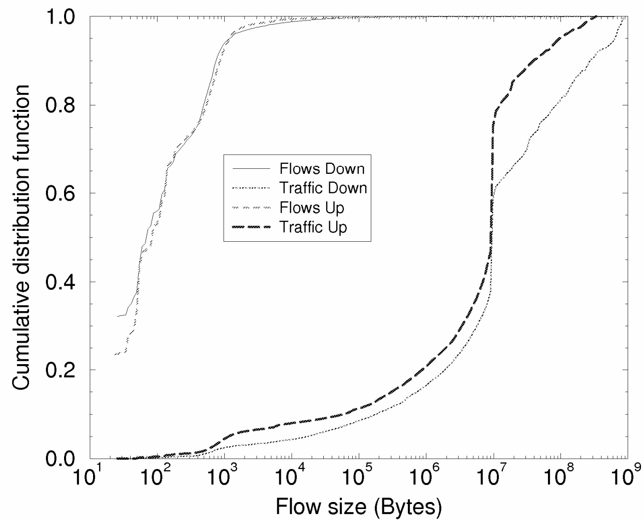

Figure 5: Cumulative distribution functions of flow sizes, vs. their number and their traffic

3.2 Loss rate and connection durations

Figure 6 presents the cumulative distribution function of the connections' duration in seconds according to some desequencement rates given in the legend. We do not put all the curves to keep a readable figure. Connections without packet loss are not presented. We observe that most of the connections last more than 1000 seconds as long as the loss rate remains lower than 10^{-3}. For higher loss rates, the curves are gradually shifted to the left. The connections are then shorter on the higher desequencement rates. We note that with the loss rates larger than 10^{-2}, the curves remain close, between 30 and 100 seconds. We deduce from these figures a clear correlation between the distribution of the duration and the loss rate of TCP connections. There may be actually several reasons to explain the reduction of the durations in case of bad performance and, more generally, the relation between the loss rate and connections' duration. Connections may be interrupted by applications or users, but users may also hesitate to launch big transfers, and applications may control their throughput by themselves. This is particularly the case for some peer-to-peer applications. They reduce their traffic in case of bad performance and may obtain lower loss rates than, for example, short web requests launched often in parallel to download the elements of a web page simultaneously. We evaluate the impact of the losses on the interruption of connections in subsection 3.4, and differentiate these impacts according to the application in Section 4.

In order to detect performance problems and unsatisfied users, we may then define two thresholds on the loss rate. These thresholds characterize three behaviours:

- Long connections with loss rates lower than 10^{-3},
- Short connections with loss rates greater than 10^{-2},
- Intermediate behaviour with loss rates between 10^{-3} and 10^{-2}.

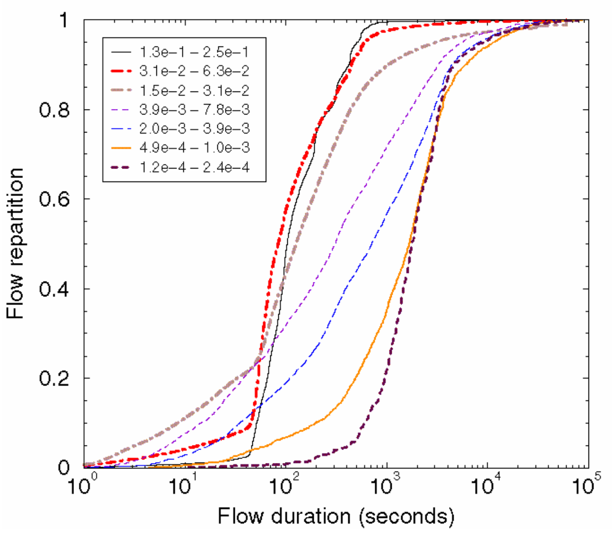

Figure 6: Distribution of flow durations for various desequencement rates

Looking at the statistics of connections durations according to the desequencement rate presented on Figure 7, we note that previous thresholds are less visible than on the distributions. This figure shows the mean flow duration and its standard deviation in seconds and their ratio, the coefficient of variation. We only observe a slight decrease of the durations for loss rates larger than 10^{-3}, then a little growth between 10^{-2} and $4{\cdot}10^{-2}$. This increase may be caused by retransmission delays, usually much greater than transmission delays due to the granularity of retransmission timers.

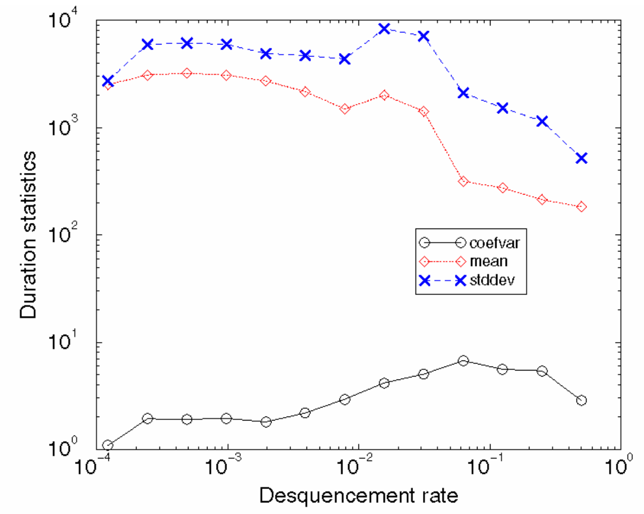

Figure 7: Statistics of flow duration versus the desequencement rate

For loss rates higher than 5%, the average durations are reduced to less than 200 seconds (which means by a factor of 10 compared to small loss rates). We observe also that the coefficient of variation increases with the loss rate. The coefficient of variation is indeed lower than 2 for loss rates lower than 10^{-3} and it increases up to 10 for loss rates of several percents. This may be explained by the fact that, for instance, in case of overload, some TCP connections are severely impacted and encounter a high loss rate, whereas some other connections do not see the overload, as we observe in [Brown and Collange 03]. Thus the dispersion of durations increases with the loss rate. For very high loss rates the decrease in the coefficient of variation may indicate that less connections "escape" to bad performance.

3.3 Loss rate and connection sizes

We now consider the correlations between connections' sizes, and retransmissions and desequencements. We use equally the size or the volume of connections to talk about their number of (upward or downward) bytes. Figure 8 presents the mean volume of connections in both directions (up and down) as the functions of the losses due to desequenced packets and to the retransmitted ones. As for duration, the size of connections is clearly related to their loss rates. We first notice that the decreasing rate is much larger than one. This excludes then the sole impact of the "size effect", i.e. the fact that a connection of 10 packets losing packets observes at least a 10% loss rate. We note a similar threshold of $4{\cdot}10^{-3}$ when the size of connections starts to decrease; it was 10^{-3} for durations. We see continuous and more important decrease than for durations, of many orders of magnitude. The upper threshold, when the size stops decreasing, is around 10% for retransmissions. There is no upper limit for desequencements; the curves are decreasing until the highest loss rates. Moreover, the increase around 1% is not observed on desequencements' curves. This confirms that the recovery observed on the duration analysis should be due to retransmission delays.

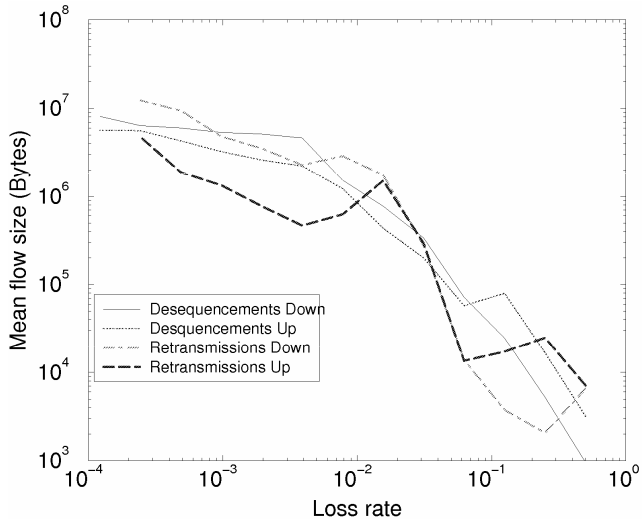

Figure 8: Mean flow size vs. loss rate

We point out from the comparison of Figure 8 and Figure 5 a possible influence of the retransmission delays in durations. Due to the large granularity of the TCP retransmission timer, the retransmission delays of the lost packets represent a significant part in the total duration of connections in case of high loss rates [Padhye et al. 98]. Thus, retransmission delays may distort the statistical results we obtain with durations and hide the possible impact of durations on the users' behaviour. The size of connections is not so impacted by the retransmitted packets. In order to detect a possible influence of the end-to-end QoS on the users' behaviour, and to bring to light performance problems, it is consequently more relevant to consider the correlations of the loss rate with transfer sizes rather than with durations.

3.4 Interruptions of connections

We have noted in the former subsections that the duration and the mean size of connections decrease when the loss rate increases. Therefore we analyze now the impact of packet losses on the proportion of interrupted connections. We use the TCP flag of their last packet in order to detect interrupted TCP connections. A host sends a packet with the R (Reset) flag [Peuhkuri 02] to interrupt a TCP connection. We note that terminating a TCP connection with a Reset flag may also sometimes be a manner of cleaning up the memory resources at server side. These connection interruptions may be caused by the user, but also at any protocol level, for example by a low level disconnection. A user may also stop a transfer when he gets sufficient information, independently of the transfer performance. So the variations of the interruption rate according to the loss give only an indication on possible users' feelings, and not their actual behaviour.

In Figure 9, we present the proportion of the last flags according to the desequencement rate. Each curve shows the proportion of connections ending by a given TCP flag in the last packet, according to the loss rate. The legend gives the last flag associated to each curve. Some curves are too close to the x-axis to be visible. For a desequencement rate of 10^{-4}, for instance, about 80% of the connections are terminated by an acknowledgment

We observe larger proportion of interrupted connections terminated by a Reset flag, even with the lowest loss rates. 7% of the connections are interrupted for a packet loss rate of 10^{-4}. This may be due to some applications systematically ending their TCP connections with a Reset flag. We further notice that the proportion of Reset flags increases with the loss rate and reaches approximately 40% of the terminations of connections for a 1% loss rate. This proportion remains roughly constant up to the loss rate of 20%. Beyond that threshold, the proportion of interrupted connections decreases. In the same time, the proportion of connections finishing with the S (Synchronization) flag increases quickly. This indicates an increasing proportion of connections that can not be established due to actual losses in the network nodes or anomalies, such as bad IP address, disconnected host and so on.

The increasing number of interrupted connections (for loss rates lower than 1%) explains the decrease in the flow sizes observed in the former subsection. In that sense, it could be also interesting to show the correlation between interruptions of connections and connection sizes.

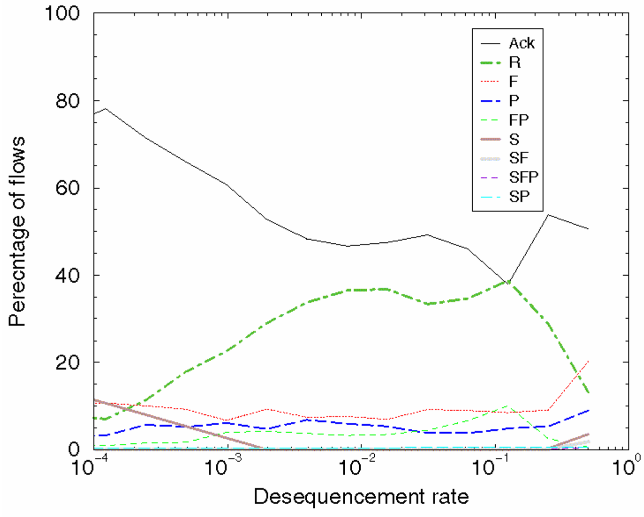

Figure 9: Distribution of last TCP flags according to the loss rate

For loss rates larger than 1%, the proportion of interrupted connections does not evolve any more. So the reduction in the mean connection size might be explained by self-censorship of the users: users do not launch or postpone long transfers that they would have launched with better performance. But it is not possible to verify this assumption on our passive measurements. It can only be checked during experiments on the test-bed, asking the users about their actual behaviour.

3.5 Arrival rate

We study now the impact of the loss rate on the inter-arrival times between consecutive connections per customer. Different behaviours may occur a priori in case of bad performance. The users may be discouraged, postpone or do not launch at all some of their transfers. In this case we should observe larger inter-arrival times. On the contrary, they may be impatient, abort some transfers in progress and restart them quickly. In this second case, the interval between transfers should then be reduced.

On Figure 10 we present the average and the standard deviation of the transfer inter-arrivals per user, according to the desequencement rate. We observe that these inter-arrivals increase with the loss rate. According to subsection 3.2, the durations of the connections decrease in such case. So the silences between transfers grow as much. From these observations, we might deduce rather discouragement of users with growing loss rates. An exponential backoff behaviour at the application level could also explain this phenomenon. As with the interruptions of connections, the actual cause can only be obtained by asking a panel of users about their real behaviour. Besides, we do not observe any clear threshold in the growth of the inter-arrivals, which appears from the smallest loss rates.

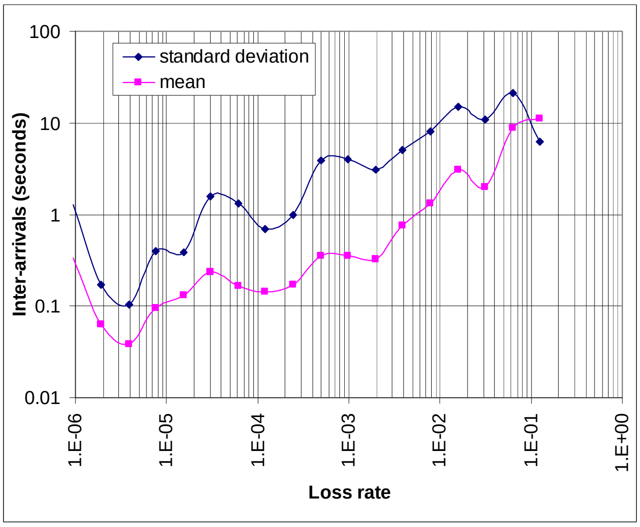

Figure 10: Inter-arrivals statistics according to the desequencement rate

4 Sensitivity of some applications to the losses

We may assume that interactive applications are more sensitive to losses than background transfers. If the users do not wait for an immediate answer they should be less reactive to bad end-to-end performance. Moreover, some applications, such as peer-to-peer applications, generate large traffic volumes and control their throughput. The TCP flow control, whose dynamic congestion window may cause packet losses, is then bypassed. Such applications may decrease the observed loss rates of big transfers. Conversely, other applications generating short and frequent simultaneous transfers such as HTTP, may observe higher loss rates. The observations of the previous sections might only be due to the differences between the traffic characteristics of the applications and not due to the actual influence of the losses, if for example applications using small transfers observe for some reasons higher loss rate. We show in this section that the losses have an actual impact, no matter of the application, and we try to quantify this impact.

4.1 Approximations of flow sizes using classical laws

Following subsection 3.3 we propose here to approximate the mean flow size according to the loss rate using some classical laws. We have considered various regressions: linear, logarithmic, exponential and power law. We observe that the latter presents the best coefficient of correlation. This correlation, already good, is still better if we consider loss rates larger than 10^{-3}. As we already noted, lower loss rates are much less correlated with the connection sizes due to the loss rate inherent to the TCP oscillating behaviour. We find then for the whole traffic the following regressions, with very good (close to 1) coefficients of correlation r:

- Downward size: $y = 3.6x^{-1.26}$, r = -0.99;

- Upward size: $y = 7.3x^{-1.10}$, r = -0.99.

We can also compute in the same way the regressions of the inter-arrivals between connections according to the loss rate (see Figure 12):

- Inter-arrivals: $y = 1.3x^{0.45}$, r = 0.95.

The powers of these regressions characterize the decrease of the mean size or the increase of the average inter-arrival, when the loss rate increases. It can be used to quantify the correlations between the loss rate and these traffic features, answering the question how fast the mean size decreases when the loss rate increases. If we use now similar regressions for an application, this method may be used to describe and compare the sensitivity of its traffic characteristics to the loss rate. The higher the exponent is, the smaller the mean flow size is.

4.2 Impact of the losses on the applications

According to the observations in Section 3, we consider the impact of the loss rate on the downward size of connections and on their inter-arrivals, per application and per user. As the power regression appears to be the most adapted law to the whole traffic, we only consider this regression to compare the sensitivity of the applications to the losses. We must note that this sensitivity does not only depend on the behaviour of the users, but also on the implementation of the applications.

Application	% connections	% bytes	Impact on sizes		on inter-arrivals	
			Slope	r	Slope	r
telnet	0.05	0.17	-2.78	-0.97		
edonkey	42.2	67.9	-1.70	-0.96	0.45	0.97
irc	1.7	1.0	-1.55	-0.92	0.20	0.90
citrix	0.04	0.23	-1.52	-0.96		
smtp	0.1	1.2	- 0.98	-0.98	0.75	0.98
http	1.9	5.8	-0.91	0.96	0.29	0.95
pop3	0.50	0.3	-0.87	-0.85		
netbios	0.15	0.27	-0.75	-0.69		
ftp	0. 20	0.56	-0.75	-0.77		
bittorrent	0.25	0.62	-0.57	-0.72		

Table 1: Sensibility of the applications to the packet losses

Table 1 gives the powers and the coefficients of correlation for the most frequent applications, characterizing the dependence of the size and of the inter-arrivals on the loss rate. These values are, of course, absolutely not general. They only correspond to measurements over one week in 2006 on a specific ADSL Point of Presence. They perhaps do not depend too much on the different points of presence of an ISP in a given country. But they certainly depend on the country, as well as on the application usage, or on the users' behaviour etc. Moreover, these values evolve over time for the same reasons and also due to new releases of protocols and applications. Hence, the ISP must frequently refresh these measurements. The slopes of some applications' inter-arrivals could not be calculated due to insufficient number of measurements. Some applications use only few TCP connections per day, for a given user. We also present the proportion of connections and bytes of each application. The applications are classified according to the influence of the losses on their downward transfer size.

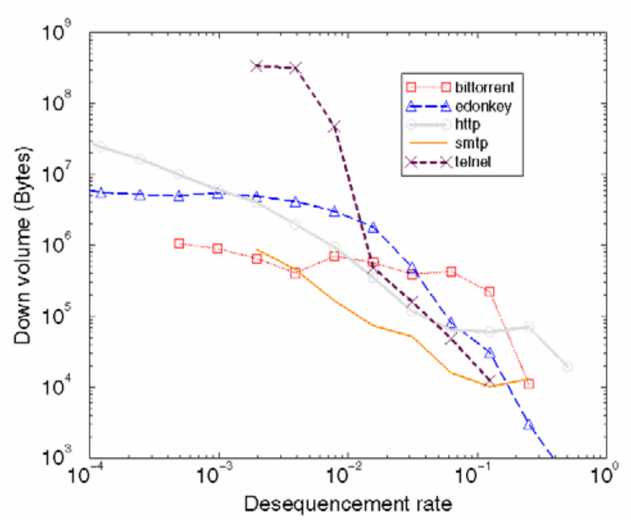

Figure 11: Mean volume by application according to the desequencement

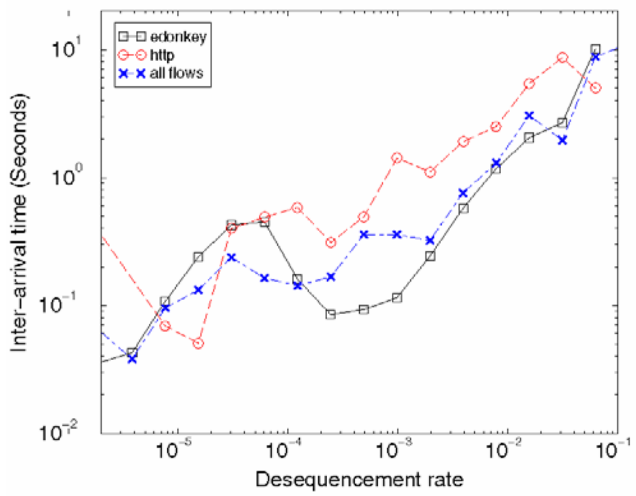

Figure 12: Inter-arrival time by application versus the desequencement rate

At first we note a not-null slope for all of these applications. This proves that the impact of packet losses observed in the previous sections for the total traffic is not an artefact due to different traffic characteristics of the applications. Secondly we observe relatively different slopes according to the application. They are particularly high for usually interactive applications (i.e. telnet, irc, citrix). These applications are, as expected, more sensitive to the losses. The position of eDonkey is more surprising. We present for few applications, on Figure 11 and Figure 12 respectively, the downward volume and the inter-arrival durations according to the loss rate. On Figure 11 we visualize more clearly the differentiated sensitivity of applications to the loss.

5 Conclusion

Based on a detailed and large set of real and operational traces, we proved the correlations between some traffic features and the network performance. The higher the loss rate, the shorter and smaller the TCP connections. We observed also that more connections are interrupted. We noted that these correlations may be, directly or through different protocol levels, due to the users' behaviour, according to their feelings on the subjective performance. But these assumptions still need to be checked on a panel of users, on controlled experiments. Our passive analysis is not sufficient to get the actual behaviour and feelings of the users. So now we are planning to perform the experiments on few interactive applications, with panels of users on test-beds. The objective of these experiments will be to compare the correlations observed in this paper with the users' answers on their actual perception and behaviour. Our aim in this paper was to propose a new method to describe the correlations between the Quality of Experience and some network level performance criteria. As stated in the introduction, we think that the three methods we described are complementary and should be used in parallel.

We also proposed a method to compare the different sensitivities of the applications to the performance. This may help to improve QoE prediction, taking into account the applications' usage forecasts. This method can also be used to prioritize some applications, but this decision would depend also on their crucial importance for the users. As the QoE depends certainly not only on the loss rate, we are now studying the correlations between traffic characteristics and other performance criteria, such as round-trip time, throughput etc.

References

[Brown and Collange 03] Brown, P., Collange, D.: "Sharing resources in overload with TCP"; ARC TCP, ARC Workshop on TCP, Paris (France), November 2003. Extended release: France Télécom Technical Note NT 8555.

[Brosh et al. 05] Brosh, E., Lubetzky-Sharon, G., Shavitt, Y.: "Spatial-Temporal Analysis of passive TCP Measurements"; IEEE Infocom, Miami, March 2005.

[Collange 07] Collange, D.: "Passive estimation of network utility functions"; 2nd workshop on Socio-Economic Aspects of Next Generation Internet, Santander (Spain), 28-29 June 2007.

[Collange et al. 05] Collange, D., Costeux, J.L., Plissonneau, L.: "Detection and Localization of Performance Problems of TCP Connections on ADSL"; EuroNGI Workshop on QoS and Traffic Control, Paris, December 2005.

[Collange and Costeux 07] Collange, D., Costeux, J.L.: "Correlation of Packet Losses with Some Traffic Characteristics"; PAM'2007, conference on Passive and Active Measurements, Louvain (Belgium), April 2007.

[Costeux and Guyard 06] Costeux, J.L., Guyard, F.: "Detection and comparison of RTP and Skype Traffic and Performance"; Globecom 2006, 49th IEEE Global Telecommunications Conference, November 27, December 1st 2006, San Francisco, USA.

[Fiedler et al. 04] Fiedler, M. (ed.): "State-of-the-art with regards to user perceived Quality of Service and quality feedback"; EuroNGI Deliverable D.WP.JRA.6.1.1, May 2004.

[Fiedler et al. 05] Fiedler, M., Chevul, S., Radtke, O., Tutschku, K., Binzenhöfer, A.: "The Network Utility Function: A practicable concept for assessing network impact on distributed services"; pp. 1465-1474 in proc. 19th International Teletraffic Congress (ITC19), Beijing, China, 2005.

[Guillemin et al. 03] Guillemin, F., Robert, P., Zwart, B.: "Heavy tailed M/G/1-PS tails with impatience and admission control in packet networks"; in IEEE Infocom, San Francisco, April 2003.

[ITU 03] ITU-T Recommendation P.800: "Mean Opinion Score (MOS) Terminology"; 2003

[ITU 05] ITU-T pre-published Recommendation G1030: "Estimating end-to-end performance in IP networks for data applications"; 2005.

[Jaiswal et al. 03] Jaiswal, S., Iannaccone, G., Diot, C., Kurose, J., Towsley, D.: "Measurement and Classification of Out-of-Sequence Packets in a Tier-1 IP backbone"; IEEE Infocom, 2003.

[Jacobson et al. 89] Jacobson, V., Leres, C., and McCanne, S.: "TCPdump"; Lawrence Berkeley Laboratory, June 1989.

[Nokia 04] Nokia White Paper: "Quality of Experience (QoE) of mobile services: can it be measured and improved"; 2004.

[Massoulié and Roberts 99] Massoulié, L., Roberts, J.: "Bandwidth sharing: objectives and algorithms"; INFOCOM'99, Conference on Computer Communications, vol.3, pp.1395-403.

[Padhye et al. 98] Padhye, J., Firoiu, V., Towsley, D., Kurose J.: "Modelling TCP throughput: A simple model and its empirical validation"; ACM Sigcomm, Vancouver (Canada), September 1998, pp. 303-314.

[Peuhkuri 02] Peuhkuri, M.: "Internet Traffic Measurements - Aims, Methodology and Discoveries"; Licenciate Thesis, Helsinki University of Technology, Finland, 2002.

[Rewaskar et al. 06] Rewaskar, S., Kaur, J., Smith, F. D.: "Accuracy of Probing Techniques in Estimating TCP Loss Rates"; ACM Sigcomm, Pisa, September 2006.

[Roberts and Bonald 03] Roberts J., and Bonald, T.: "Congestion at flow level and the impact on user behaviour"; Computer Networks 42, 2003, pp 521-536.

[Rossi 02] Rossi, D.: "At the Network Edge: Sniff, Analyze, Act"; Ph.D. thesis, Politecnico di Torino, Departimento di Elettronica, 2002.

[Yang and Veciana 01] Yang S., Veciana, G.: "Bandwidth sharing: the role of user impatience"; IEEE Globecom, 2001, pp. 2258-2262.

[Zhao et al. 06] Zhao, Y., Chen, Y., Bindel, D.: "Towards Unbiased End-to-End Network Diagnosis"; ACM Sigcomm, Pisa, September 2006.