Internet traffic characterization – An analysis of traffic oscillations

Philippe OWEZARSKI, Nicolas LARRIEU

LAAS - CNRS
7, avenue du Colonel ROCHE
31077 TOULOUSE Cedex 4 FRANCE
{owe, nlarrieu}@laas.fr

Abstract. Internet traffic has been changing a lot since few years in particular with the arrival of new P2P applications for exchanging audio files or movies and nowadays the knowledge we have on it is quite limited. Especially, new applications and new traffic are creating a lot of troubles and performance issues. Based on some traffic traces captured in the framework of the METROPOLIS network monitoring project, this paper exhibits the highly oscillating nature of Internet traffic, thus explaining why it is almost impossible nowadays to guarantee a stable QoS in the Internet, and also that such oscillations provoke a huge decrease of the global network QoS and performance. This paper then demonstrates that traffic oscillations can be characterized by the Hurst (LRD) parameter. In particular, this demonstration relies on a comparative study of Internet traffic depending on the transport protocol used to generate it. It is then shown that using TFRC – a congestion control mechanism whose purpose deals with providing smooth sending rates for stream oriented applications – instead of TCP, makes traffic oscillations and LRD almost disappear. This result, i.e. limiting as much as possible the oscillations of traffic sources in the Internet, then gives research directions for future Internet protocols and architectures.

Keywords. Internet monitoring, traffic characterization, LRD, TFRC

1 Introduction

The Internet is on the way of becoming the universal communication network for all kinds of information, from the simple transfer of binary computer data to the transmission of voice, video, or interactive information in real time. It has then to integrate new services suited to new applications. In addition, the Internet is rapidly growing in size (number of computers connected, number of users, etc.), and in complexity, in particular because of the need of new advanced services, and the necessity to optimize the use of communication resources to improve the Quality of Service (QoS) provided to users. In fact, the Internet has to evolve from a single best effort service to a multi-services network.

Such evolution is not that easy due to the complexity of the Internet and all its network interconnections, with their resource heterogeneity in terms of technologies but also in terms of provisioning, and of course with the traffic characteristics. Indeed, all new applications with various and changing requirements introduce in Internet traffic many characteristics that are very far from common beliefs. The increase of the Internet complexity leads to some difficulties in understanding how all the Internet components and protocols interact. In the same way, the control we can have on the Internet – or even on small parts of the Internet – is quite limited. The evolution of the Internet is then strongly related to a good knowledge and understanding of traffic characteristics that will indicate the kind of mechanisms to deploy to perfectly match user requirements and network constraints. Consequently, the development of monitoring-based tools and technologies to collect Internet traffic information, and methodologies to analyze their characteristics is currently an important topic for network engineering and research.

This paper then deals with presenting the first results of the METROPOLIS project on traffic characterization and analysis. METROPOLIS is a French national project granted and funded by the French Network for Research in Telecommunications (RNRT[1]) whose main goal deals with issuing new network monitoring and analysis methodologies. The network under consideration is RENATER, the French network for education and research. The traces are captured using passive measurement tools that collect the TCP/IP headers of every packet with a very accurate GPS timestamp. The capture system is based on the DAG card [2].

In particular, the first traffic traces captured and analyzed exhibited a very important and unattended oscillating nature, that then makes us understand why it is so difficult to guarantee QoS in the Internet. But such a result is useless if we are not able to analyze and explain its causes. This paper then deals with analyzing the oscillating nature of Internet traffic, as well as trying to model it qualitatively and quantitatively. For this purpose, this paper then explains why traffic oscillates. This explanation relies on some lacks of TCP that is not perfectly suited for the transmission of the more and more frequent long flows (as movies or music files) on high speed links (section 2). This paper then presents an analysis of the oscillations features as amplitude and range. Especially, it is shown in section 2 that Long Range Dependence (LRD) can model the oscillations of Internet traffic, and in particular, their range (what is the more important aspect of oscillations related to networking issues). Section 3 then proposes to illustrate on a case study the explanation proposed in section 2. The principle of this example based demonstration consists in studying and analyzing Internet traffic if TCP is replaced by a smooth transport protocol. The TCP Friendly Rate Control (TFRC) has been designed for providing a smooth sending rate for stream oriented applications. Both traffics, i.e. using TCP vs. TFRC, are then comparatively analyzed, thus showing the impact of TFRC on oscillations. Based on this analysis result, section 4 concludes this paper by giving some research directions for future network protocols and architectures able to improve Internet services.

---

[1] RNRT: Réseau National de Recherche en Télécoms

2 Traffic oscillation issues and elephant flows

This first section is then presenting the characterization and analysis results we got after analyzing the traffic traces that have been captured on RENATER network. To well understand the new traffic characteristics, it is first required to analyze the evolution of the Internet in terms of usages.

2.1 Evolution of traffic characteristics

The evolution of Internet traffic these last years has been marked by the huge increase of P2P traffic (Kaaza, e-donkey, ...), and now, on some links of the RENATER network, it can represent the same proportion than HTTP traffic (Figure 1). Such a result is quite impressive because, in an academic network as RENATER, students, teachers and researchers are not supposed to download music or movies. And, in fact, the amount of P2P traffic in RENATER is pretty low compared to the results observed on the commercial network of France Télécom[2], especially on the ADSL POP were P2P traffic can grow up to 70 % – and sometimes more!

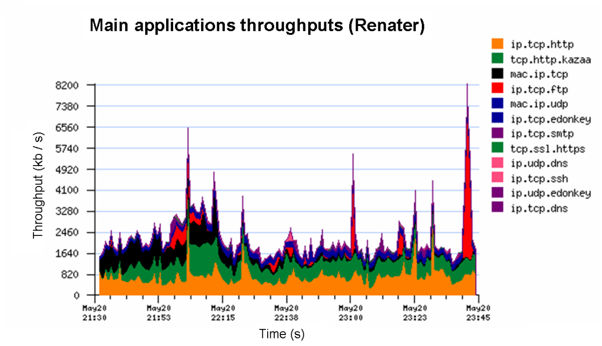

Fig. 1. Traffic distribution on the RENATER network in May 2003

Such an increase of P2P has necessarily an impact on traffic characteristics. In particular, because of the nature of file exchanged – mostly music and movies – that are very long compared to web traffic that was the dominant traffic in the Internet few years ago. In fact, the increase of P2P traffic, in addition of the classical traffic, makes the traffic have the following characteristics:

---

[2] France Télécom R&D is part of the METROPOLIS project, but the results got on the France Télécom network are not public and will not be more discussed in this paper.

- There are always thousands of mice[3] in Internet traffic (because of the web, as well as P2P controls);
- But there are also a large number of elephants.

So, one of the main consequence of the evolution in terms of applications and usage is related to the flow size distribution changes. Figure 2 represents the flow size distribution between 2000 and nowadays. The exponential function (in black) is taken as a reference because the exponential distribution is closely related to the Poisson model that is most of the time used as the reference model for Internet traffic in simulations or for performance evaluation. We can see on this figure that between year 2000 and nowadays the proportion of very long flows has increased in an important way. If in 2000, flow size distribution was almost exponential, this is completely wrong nowadays. Current distribution is very heavy tailed, and this distribution is then very far from the exponential distribution traditionally considered.

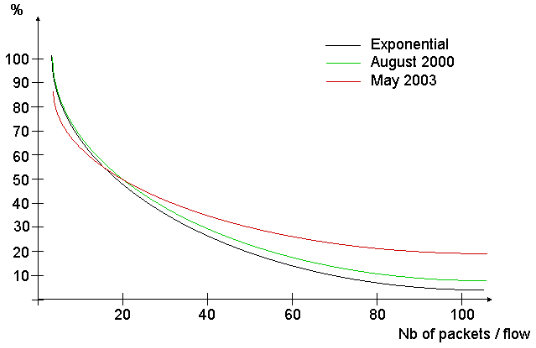

Fig. 2. Flow size distribution evolution between 2000 and 2003

2.2 Traffic Long Range Dependence and related issues

This increase of the proportion of P2P elephants hugely impacts traffic profile. Figure 3 illustrates it in current traffic. It shows the difference between the actual Internet traffic and Poisson traffic. These two traffics are observed with different granularities (0.01 s, 0.1 s and 1s), and it appears that Internet traffic does not

---

[3] "Mouse" is a term used to designate a small flow, i.e. a flow that does not last enough to exit from the slow start phase of TCP. At the opposite, very long flows are called elephants.

smooth as fast as Poisson traffic when increasing observation granularity. The analysis demonstrated that this result is completely due to elephants. In fact, the transmission of elephants creates in the traffic the arrival of a large wave of data that has the particularity of lasting for a long time – more than 1 second – while web flows are generally transmitted in less than one second on the current Internet. That is why we have this difference between Poisson and real traffic: the nature of oscillations between the two traffics changes, with oscillations in actual current traffic that are more persistent.

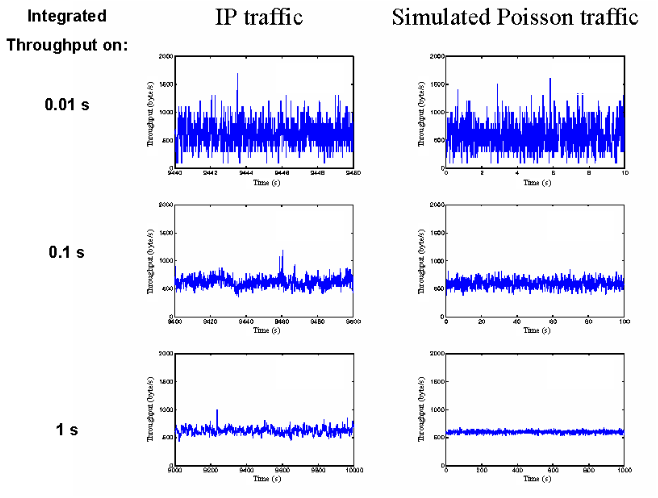

Fig. 3. Comparison between oscillations of Internet and Poisson traffics

In addition, as TCP connections used for transmitting larger flows are longer, the dependence that exists between packets of the same connection propagates on longer ranges. This phenomenon is usually called Long Range Dependence or long memory. It has several causes, and in particular congestion control mechanisms deployed in the Internet, especially the ones of TCP, this protocol being the dominant protocol in the Internet [9]. Among all the TCP mechanisms, it is obvious that its closed control loop introduces dependence, as acknowledgements depend on the arrival of a packet, and the sending of all the following packets of the connection depends on this acknowledgement. In the same way, the two TCP mechanisms – slow-start and congestion avoidance – introduce some dependence between packets of different congestion control windows. By generalizing these observations, it is obvious that all packets of a TCP connection are dependent the ones from the other. In addition, with the increase of the Internet link capacities that allows the transmission of longer and longer flows, it is obvious that the range of the LRD phenomenon increases. That is why the persistence of the Internet traffic oscillations measured, even with a coarse granularity, is so high. Indeed, because of the TCP dependence phenomenon propagating in the traffic thanks to flows (connections), the increase of flow size also makes the dependence range increase and propagates on very long ranges. An oscillation at time t then provokes other oscillations at other time being potentially very far from t. A (short term) congestion due to a huge oscillation of a connection can then continue to have some repeats several hours later (in the case of a movie download for instance), i.e. this flow will continue to propose to the network some traffic peaks directly dependent from this first oscillation, and can create some new short term congestions. Moreover, it is clear that elephants, because of their long life in the network, and because of the large capacities of networks – most of the time over-provisioned – have time to reach high values of the congestion control window (CWND). Thus, a loss induces a huge decrease, followed by a huge increase of the throughput of the flow. The increase of flow size then favors high amplitude oscillations, dependent on very long ranges. Of course, oscillations are very damaging for the global utilization of network resources as the capacity released by a flow that experiences a loss (for example) cannot be immediately used by another flow (because of slow start for instance): this corresponds to resource waste, and introduces a decrease of the global QoS of the network. In fact, the more the traffic oscillates, the lower the performances [10].

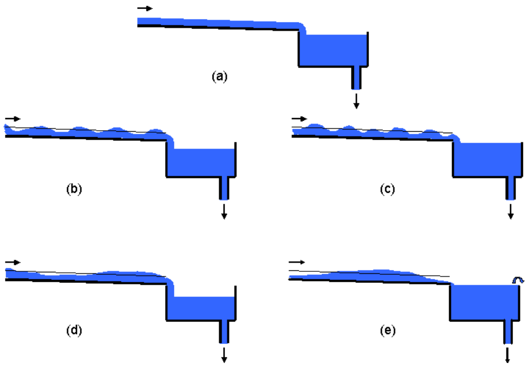

Fig. 4. Illustration of LRD issues on losses

To give a concrete view of LRD issues on traffic, Figure 4 aims at illustrating it. Figure 4.a depicts a leacky bucket as an analogy with a router, for instance, its buffer, ingress and egress links. So, when there are waves in the arriving traffic (Figure 4.b), and if the goal is to provide a good service with no extra losses and no extra delays, it is first required to over-provision the link (otherwise the traffic will be smoothed, and at least, delays will be introduced for some packets). The second characteristic appears in the buffer when a wave arrives: it makes the level of the buffer increase (Figure 4.c). This is a well known issue of networking addressed many time before, especially by [4]. But when the range of oscillations increases (Figure 4.d), and this is the case with current Internet traffic, the arrival of a persistent wave provokes a buffer overflow, thus leading to losses (Figure 4.e). As a conclusion of this practical illustration, it is important to point out that LRD in traffic induces bad performances and QoS for networks, as it is the source of congestion and losses.

Figure 5 represents the LRD evaluation of the traffic depicted in Figure 3. This figure has been produced using the LDestimate tool [1] which estimates the LRD that appears in Internet traffic at all scales[4]. The principle of this tool relies on a wavelet decomposition of traffic time series, what then allows users to have a graphical representation of the dependence laws at all time scales, i.e. the variability of oscillations depending on the observation range. In figure 5, we can note a "bi-scaling" phenomenon (two lines in a log-log scale on Figure 5 with an elbow around octave 8) which shows a difference in the LRD level between short and long time scales for the traffic exchanged. For short scale (octave < 8), representing the dependence between close packets (i.e. packets whose sending time are not very far from each other), the dependence is quite limited. Such dependence is the one that can exist for packets belonging to the same congestion window and that are then very close from each other. On the other side, for long time scales (octave > 8) LRD can be very high. For octaves 8 to 12, that correspond for instance to the dependence between packets of consecutive congestion windows, the dependence is higher. This can be explained by the closed loop structure of TCP congestion control mechanism in which the sending of one packet of a congestion control window depends on the receiving of the acknowledgement of one packet of the previous congestion control window. Of course, this phenomenon exists for consecutive congestion window, but also for all congestion windows of the same flow. This means, that the presence in the traffic of very long flows introduces very long scale dependence phenomenon, as depicted on figure 5 for very large octaves. What comes out from this LRD analysis is a bad utilization of resources as TCP is not suited for the transmission of long flows on high speed networks, with an increase of LRD and oscillations and thus a decrease of the QoS.

3 A case study illustrating the relations between oscillations and LRD in Internet traffic

The last section showed two phenomena of Internet traffic: oscillations and LRD. These observations and analysis, combined with the results published in the literature [5] [6] [7] [9] [10], etc. on traffic characterization, make us think that LRD is a good way to characterize traffic oscillation phenomenon, in particular for their "range" feature. Up to our knowledge, the problem of oscillation range characterization has not been addressed before in the existing literature. Thus, the following experiment aims, on a specific case, at showing the close

---

[4] Note that the Hurst factor H that is the parameter fully characterizing a self-similar process – and Internet traffic if often said to be self-similar [9] – can be obtained directly depending on the slope of the LRD curve.

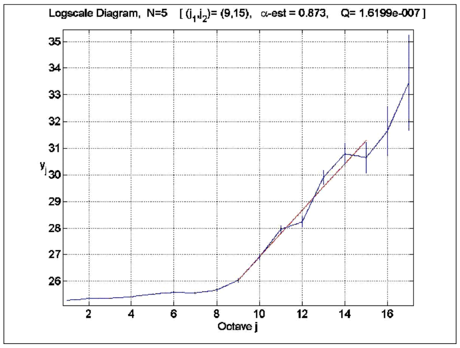

Fig. 5. LRD evaluation for the traffic of an Internet access link

relationship that exists between oscillations range and LRD. For this purpose, the proposed experiment compares, based on NS-2 simulations, the actual traffic with the same replayed traffic for which the TCP transmission protocol has been replaced by TFRC [5] [8]. The goal of TFRC is to provide traffic sources that have a more regular and smooth sending rate compared to TCP[5]. This case study then aims at showing that when using TFRC, i.e. when generating regular and smooth traffic, LRD is quite reduced compared to the TCP case.

3.1 TFRC principles

The TFRC rate control computes once by round trip time (RTT), the sending rate for each source. This computation depends on the loss event rate measured by the receiver [3] according to equation 1:

$$X = \frac{s}{R * \sqrt{2 * b * \frac{p}{3}} + (t_{RTO} * (3 * \sqrt{3 * b * \frac{p}{8}}) * p * (1 + 32 * p^2)))} \tag{1}$$

where:

- X is the transmit rate in bytes/second,
- s is the packet size in bytes,
- R is the round trip time in second,
- p is the loss event rate (between 0 and 1.0), of the number of loss events as a fraction of the number of packets transmitted,
- t_{RTO} is the TCP retransmission timeout value in second,
- b is the number of packets acknowledged by a single TCP acknowledgement.

In TFRC, a loss event is considered if at least one loss appears in a RTT. This means that several losses appearing in the same RTT are considered as a

---

[5] Initially, TFRC has been designed for stream oriented applications, for instance the ones transmitting audio or video flows in real time.

single loss event. Doing so, the loss dependence model of the Internet is broken since most dependent losses are grouped in a same loss event (related to short term congestions). Thus, the recovery will be eased and more efficient compare to what TCP can do: it is well known that TCP is not very efficient to recover from several dependent losses or losses in sequence.

3.2 Experiment description

This experiment aims at providing a comparative evaluation of the global traffic characteristics if elephants use TCP or TFRC as the transmission protocol. This experiment aims to provide values in a realistic environment. For that, the experiment relies on the use of traffic traces grabbed thanks to passive monitoring tools as the DAG [2] equipments. Thus, traffic flows identified in the original traffic trace are replayed in NS-2 with the same relative starting date and respecting all the others characteristics (packet sizes, flow sizes, etc.). On the other side, the simulation environment has been built in order to allow a coherent shaping of packets related to what happened in reality. Thus, queues and links in the simulation environment are provisioned in a way that allows the enforcement of the loss rate and model observed in the actual traffic. Delays on every link are also selected in order to respect the real RTT measured on real flows. Finally, traffic sources are put at some places that allow the simulation environment to recreate the same contentions between flows than the ones that were existing with real traffic. For more information about the re-simulation method, reader can refer to [8]. In this simulation environment, our experiment consists in transmitting elephant flows using TFRC, while other flows use the classical TCP New Reno. Note that only elephants flows are replayed using TFRC. Indeed, we showed that elephants are responsible of the long range oscillations, and in addition, elephant traffic represents the larger part of Internet traffic (always more than 60 % of the total amount of traffic). The transmission of short flows (mice) do not create any troubles. The following then proposes a comparative study between the traces where elephants are generated using TCP vs. TFRC.

Given the goal of this comparative study that aims at analyzing the impact of TFRC on the oscillating nature of the traffic, the parameters that are going to be evaluated are classical traffic throughput parameters, but also traffic statistical parameters as the LRD (as justified in section 3.) and some parameters related to variability. For that, we used the Stability Coefficient (SC), that is defined as the following ratio:

$$\text{Stability Coefficient (SC)} = \frac{\text{exchanged average traffic}}{\text{exchanged traffic standard-deviation } (\sigma)} \tag{2}$$

3.3 TFRC impact on oscillations

Figure 6 presents the traffic in both cases, i.e. in the real and simulated (with TFRC) cases. It visually clearly appears that using TFRC for sending elephants, instead of TCP, makes global traffic much smoother, avoiding all the huge peaks that can be seen on the real traffic.

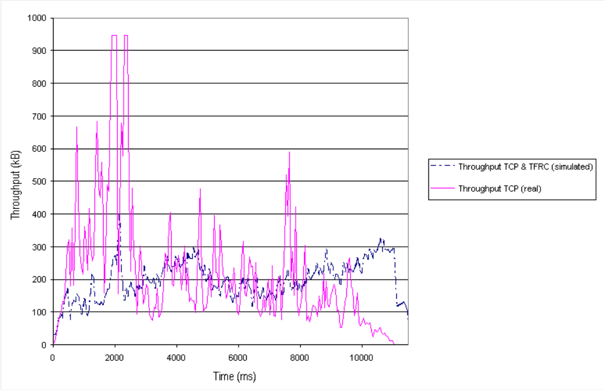

Fig. 6. Comparative evolution of the TCP and TFRC throughput

Quantitatively speaking, results are indicated in table 1. This confirms that the traffic variability in the case of real traffic (using TCP for transmitting elephants) is much more important compared to the simulated case in which elephants are transmitted using TFRC (for the standard deviation σ it has been calculated that σ(real traffic) = 157.959 ko $\ggg$ σ(simulated traffic) = 102.176 ko). In the same way, the stability coefficient is less important in the real case (SC = 0.521) than in the simulated one (SC = 0.761).

Dealing with the global throughput we got for both real and simulated traffic rather equal values (Throughput(real traffic) = 82.335 ko $\approxeq$ Throughput(simulated traffic) = 77.707 ko). This result is quite good as TFRC is not able to consume as many resources as TCP [7], and even if TFRC is less aggressive than TCP, it is able to reach almost the same performance level as TCP. This confirms the importance of stability for good performances [10].

Protocol	Average throughput (kB)	Throughput σ (kB)	SC
TCP New Reno (NR): real case	82.335	157.959	0.521
TCP NR & TFRC: simulated case	77.707	102.176	0.761

Table 1. Comparison of traffic parameters values for TCP and TFRC

Speaking about LRD in the simulated case, figure 7 shows that the bi-scaling property of the curve is strongly decreased, and that the curve has a very small slope. This means that all kinds of dependences, especially the long term ones have been drastically reduced. The values for the LRD (Hurst factor are: H(real

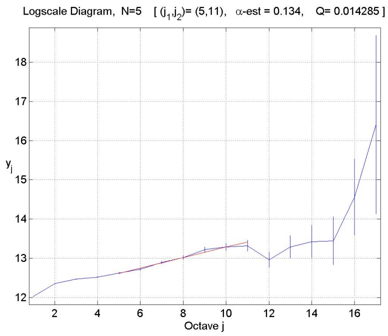

Fig. 7. LRD evaluation for simulated traffic including TFRC elephants

traffic) = 0.641 and H(Simulated traffic) = 0.194). Such result confirms two aspects of our proposal:

- TFRC helps to smooth individual flow traffic (thus providing a smoother QoS better suited for stream oriented applications) as well as the global traffic of the link;
- LRD is the right parameter to qualify and quantify all scaling laws and dependencies between oscillations.

4 Conclusion

This paper illustrated the use of traffic monitoring for network engineering and networking. Based on some traffic traces captured on the RENATER network, this paper exhibited the oscillating nature of Internet traffic (also confirmed in the quoted literature) and all the related performance issues, as well as an analysis of its causes. It then put forward the close relation that exists between traffic oscillations and LRD. Indeed, using, for transmitting most of the traffic load (i.e. elephants), a protocol that does not create oscillations (TFRC) and that breaks the dependence model between losses to recover[6], makes LRD almost disappear from traffic.

---

[6] TFRC breaks the dependence model between losses by grouping in the same loss event all losses related to the same congestion event, and by recovering all these losses at the same time. Otherwise, as in the TCP case, each loss is individually recovered, what creates traffic oscillations for each of these losses. As the losses that appear in the same RTT are certainly due to the same congestion event, they are then certainly dependent. In fact, using TCP and the single loss recovery principle, makes the dependence between losses also appear between loss recoveries, and then between traffic oscillations, because the congestions control mechanisms of TCP (that create the oscillations) react on losses. By trying to recover at the same time all dependent losses, TFRC does not propagate the dependence existing between losses on the traffic.

This analysis result is quite important as it gives us a tool for qualitatively and quantitatively characterize one of the most significant phenomena of Internet traffic, and that represents, in addition, one of the most degrading features of network performance. Especially, it gives some research directions for eliminating this bad characteristic of the traffic. For this purpose, it is then advised to look for new congestion control mechanisms providing smooth sending rates. TFRC is a candidate, but the simulation results presented in this paper also showed that TFRC is not able to reach the same global performance level as TCP, even if TCP is, in this case, not well suited (Table 1). In addition, what this paper also showed is that it should be easy to improve network performance by smoothing traffic. And this means that it is not sufficient to design new congestion control mechanisms, but also all components and mechanisms dealing with network or traffic control have to be designed for this purpose. In particular, such traffic characterization result gives good indications on how to proceed. In the same way, as oscillations have a bad impact on network performance and QoS, it is possible to propose a charging mechanism whose principle deals with charging more oscillating flows, i.e. flows that have a bad impact on the traffic [7]. Such a charging approach would aim at urging users to use congestion control, and in addition, the ones that can contribute to the smoothing of traffic.

References

1. P. Abry, D. Veitch, *Wavelet Analysis of Long Range Dependent Traffic*, Transaction on Information Theory, Vol.44, No.1, January 1998.
2. J. Cleary, S. Donnely, I. Graham, A. McGregor, M. Pearson, *Design principles for accurate passive measurement*, PAM (Passive and Active Measurements) Workshop, Hamilton, New Zealand, April 2000.
3. S. Floyd, M. Handley, J. Padhye and J. Widmer, *Equation-based congestion control for unicast applications*, In Proc. ACM SIGCOMM, 14 pages, 2000.
4. L. Kleinrock, *Queing Systems Theory*, Wiley, 1975.
5. N. Larrieu, P. Owezarski, *TFRC contribution to Internet QoS improvement*, Proceedings of the fourth COST 263 international workshop on Quality of Future Internet Services (QoFIS2003), 1 - 3, October 2003.
6. P. Olivier, N. Benameur, *Flow level IP traffic characterization* , Proc. of ITC'17 Moreira de Souza, Fonseca and de Souza e Silva (eds.), December 2001.
7. P. Owezarski and N. Larrieu, *Coherent Charging of Differentiated Services in the Internet Depending on Congestion Control Aggressiveness*, to be published in Computer Communications, special issue on "Internet Pricing and Charging: Algorithms, Technology and Applications", 2003.
8. P. Owezarski and N. Larrieu, *Trace based methodology for realistic simulations*, IEEE International Conference on Communication (ICC'2004), Paris, June 2004.
9. K. Park, G. Kim, M. Crovella, *On the relationship between file sizes, transport protocols, and self-similar network traffic*, IEEE ICNP, 1996.
10. K. Park, G. Kim and M. Crovella, *On the Effect of Traffic Self-similarity on Network Performance*, SPIE International Conference on Performance and Control of Network Systems, November, 1997.