Enhancing QoS Metrics Estimation in Multiclass Networks

Solange Rito Lima, Pedro Sousa, Paulo Carvalho
University of Minho, Department of Informatics
4710-057 Braga, Portugal
{solange,pns,pmc}@di.uminho.pt

ABSTRACT

This paper discusses the problematic of QoS monitoring, suggesting the use of on-line multipurpose active monitoring in multiclass networks as a powerful tool to efficiently assist and enhance the control of multiple service levels. To improve the simultaneous estimation of one-way QoS metrics, we propose a flexible probing source able to adjust probing patterns to the measurement requirements of each service class, exploring pattern coloring to better sense packet loss. The proof-of-concept provided shows that the proposed solution improves the estimation accuracy of multiple QoS metrics significantly, with a reduced probing overhead.

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations—*Network monitoring*; C.4 [**Performance of Systems**]: Measurement techniques

1. INTRODUCTION

Systematic network monitoring is an essential management task as it allows to: (i) keep track of ongoing Quality of Service (QoS) and network performance levels; (ii) provide feedback to traffic control mechanisms and trigger network recovery procedures; and generically (iii) support traffic engineering tasks. In this way, the research community and industry has made strong efforts to define relevant network performance metrics and to develop measurement methodologies, measurement tools and monitoring systems for their estimation and control.

In multiclass network environments, where distinct QoS profiles and service specifications need to be fulfilled, the problematic of QoS monitoring is further stressed. In these networks, QoS evaluation needs to be carried out in a per-class basis so that each class measuring requirements and behavior are met and sensed properly. The monitoring process should provide measures reflecting the real status of services' performance without introducing significant overhead or interfering with operational network traffic. Therefore, measurements have to be accurate, fast and carried out on a regular basis, while minimizing intrusion. To pursuit this objective and foster scalability of active monitoring solutions, we argue that a single probing stream should be able to capture multiple QoS metrics of a service class simultaneously, which we call multipurpose estimation.

This paper proposes the use of active multipurpose monitoring as a lightweight process to provide QoS feedback to assist service control and management tasks. In this context, a new and versatile probing source has been developed in order to improve the simultaneous estimation of one-way delay, jitter and loss-related metrics. This allows to simplify and enhance the on-line monitoring process, increasing also its ability to scale. Starting with a study of relevant characteristics of monitoring systems, the debate focuses on the problematic of QoS monitoring, identifying relevant QoS metrics and efficient strategies for their estimation, mainly from an edge-to-edge perspective (Section 2). This analysis grounds the motivation for performing on-line multipurpose monitoring in multiservice networks (Section 3). The proof-of-concept is also provided through an expressive set of simulation results (Section 4).

2. MONITORING SYSTEMS PROPERTIES

A monitoring system can follow either a centralized or distributed architecture. A centralized approach facilitates an integrated and consistent view of the network performance, but scalability problems may occur in infrastructures involving large number of monitoring nodes and significant volume of monitoring data. In distributed monitoring systems, data is collected and processed at each measurement point (MP) or, more commonly, at the receiver side of each pair of MPs, following a sender-receiver model. This latter approach is followed by most of the available freeware Software Management Tools, being widely used for on-line QoS measurement purposes, particularly on an edge-to-edge or path basis. The classic process to obtain QoS measures resorts to passive and/or active measurement methodologies, i.e., based on existing or intrusive traffic, respectively. Figure 1 illustrates these concepts.

Regarding time granularity, monitoring can be carried out off-line or on-line, i.e., based on a pos-processing or real-time data analysis. *Off-line monitoring* is more oriented to guide long-term decisions and provide a broad view of the network operation, accounting and diagnostic. *On-line monitoring* is specially oriented to provide feedback to short or medium term network management and traffic control mechanisms, i.e., the monitoring outcome is required to drive reactive

SAC'07 March 11-15, 2007, Seoul, Korea

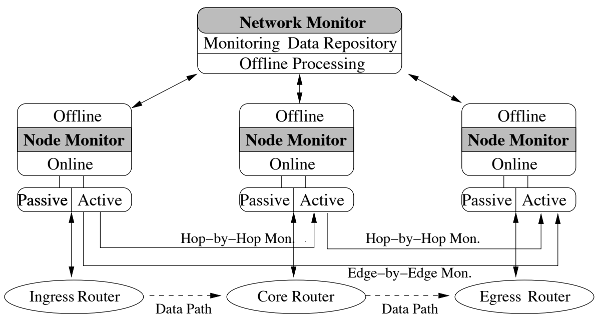

Figure 1: Example of a monitoring system

mechanisms so that traffic control decisions are not decoupled from the current network status. Currently, SLS auditing, formerly taken as an off-line task, assumes a crescent relevance as an on-line task [4], as customers are increasingly demanding in assessing the provided service levels on a near real-time basis. Performing efficient and accurate distributed on-line measurement in multiclass networks is still an open issue, requiring innovative lightweight solutions. This grounds the motivation for the present work.

2.1 Problematic of QoS monitoring

The problematic of monitoring involves the definition of adequate metrics, measurement methodologies and timing decisions, aiming at obtaining an unified and accurate view of the Internet services' performance.

2.1.1 Definition of metrics

Both ITU-T work on QoS in IP networks and the IP Performance Metrics (IPPM) WG within IETF have devoted substantial efforts to definition of metrics [5,10]. IPPM aims at developing a set of standard metrics providing unbiased quantitative measures of quality, performance and reliability of operational Internet services, proposing also measurement methodologies for the defined metrics [10]. Although adopting different terminology [5], ITU-T also identifies worst-case QoS upper bounds for common services.

One-way metrics have deserved special attention and preference over two-way metrics. In fact, due to possible asymmetric paths and/or different network resource allocation and queuing behavior in both directions, one-way measurements give more precise information and are, therefore, more useful. As regards SLS metrics, the available bandwidth, one-way packet loss, one-way loss pattern, IP packet delay variation and one-way delay are identified as the most relevant SLS metrics [6].

2.1.2 Measurement methodologies

Measurement methodologies are typically classified as active or passive. *Active measurements* resort to intrusive traffic, or probes, specifically injected in the network for measurement purposes. This type of methodology allows to emulate a wide range of measurement scenarios, providing a straightforward approach of assessing edge-to-edge QoS objectives (see Figure 2). For instance, as specific packets are injected in the network containing timestamping and sequencing data, delay and loss estimations are simplified. However, probing needs to be tightly controlled so that it does not disturb or interfere with the normal network operation. This concern is further stressed when it is carried out per traffic class.

Passive measurements use existing network traffic for metric computation. Particularly suitable for troubleshooting, passive measurements commonly resort to special-purpose devices and built-in mechanisms available in network devices. Monitoring solutions based on SNMP are representative of this type of measurement. In high-speed networks, passive measurements are a particular challenge, specially when all packets have to be accounted for, as the amount of data gathered tend to be substantial and the packet processing time very small. To deal with this, sampling techniques, more powerful hardware and new packet buffering techniques may be required [2].

Although passive techniques are usually used to monitor the performance of single nodes, they can also be applied to edge-to-edge measurements, for instance, combining hop-by-hop metrics along the network path. This allows reducing the network interference and amount of synthetic traffic, at expense of increasing processing and synchronization needs. An alternative edge-to-edge approach still within the scope of passive measurements relies on the analysis of information of real application flows (e.g., using TCP ACK or RTCP data). This approach is also referred as passive probing [3].

To take advantage of the positive aspects of both methodologies, many authors propose the use of integrated solutions, where passive and active measurements are combined to achieve more scalable monitoring systems [1, 2, 9].

2.1.3 Timing issues

Within on-line monitoring, timing issues can be twofold. At higher level, timing decisions are related to the periodicity of measurements, which depend on the measured parameters' purpose. For instance, while a time scale ranging from seconds to minutes is appropriate for AC, for active queue management (AQM) or packet scheduling the operating time scale varies from pico to milliseconds. Choosing a time scale should also consider that a small time granularity increases metrics' computation and dissemination overhead, leading eventually to excessive reaction to short-time traffic fluctuations, whereas a sparse granularity may lead to measures reflecting out-of-date network state information.

At lower level, timing decisions may require solutions to minimize or solve the problem of accurate clock synchronization between MPs in different systems. This need is notorious when measuring absolute time differences such as one-way delay, or when aggregating hop-by-hop measurements. Apart from common NTP and GPS solutions, the clock associated with Code Division Multiple Access mobile telephone network can be used as a highly accurate, synchronized distributed clock source.

3. MULTICLASS ACTIVE MONITORING

Multiclass networks pose additional challenges to on-line active monitoring. As each traffic aggregate receives a distinct treatment from either a node or domain perspective, probing needs to be carried out in a per-class basis (in-band) so that it can be adjusted to each class measuring requirements and the class behavior is correctly sensed (see Figure 2). In these networks, where the amount of measurement data increases with the number of service classes and interfaces being monitored, building scalable monitoring systems is even more relevant. Research on this topic suggests that:

the monitoring process granularity should be at aggregate level and not at microflow level; the measures transmission overhead should be minimized using event notification and statistics summarization; the amount of synthetic or intrusive traffic should be reduced [1].

Although active monitoring in multiclass networks has been matter of interest [1, 3, 12], obtaining an efficient, accurate, continuous, low-overhead and low-interfering on-line active monitoring process is still a relevant topic requiring further study. In our view, to reduce intrusion and minimize probing impact, a single probing pattern should be able to capture simultaneously multiple QoS metrics of a service class. To characterize this ability, the concept of *multipurpose* active monitoring is introduced. Investigating multipurpose probing patterns is an objective of this work.

3.1 Probing patterns

The type of probing patterns used for metric estimation varies according to the metrics to be computed and the periodicity required for their evaluation. For instance, for measuring delay and loss related parameters continuously, simple and very low rate probing patterns have been in use in real and experimental network environments (e.g., 2 or 4 packets per second (pps) [4]). For bandwidth estimation, several techniques such as Variable Packet Size, Packet Pair/Train Dispersion, Self Loading Periodic Streams and Trains of Packet Pairs have been proposed [11]. Common measurement tools based on these techniques [6] use a high volume of packets (e.g., 100-5000 pkts) or probing rate and require a significant amount of time to obtain a singleton measure (e.g., 40-100 RTTs). This slow process impairs its use for continuous estimates due to the underlying overhead, being more appropriate for sparse measurements.

For continuous measurements, common methods for collecting sample metrics use either periodic or random sampling. In *periodic sampling* measurements are made evenly spaced in time. Although being attractive for its simplicity, its eventual drawback is a possible synchronization with a periodic behavior either of the metric itself or induced by a network component. In *random sampling*, such as Poisson or geometric sampling, the samples are taken at independent, randomly generated time intervals according to a statistical distribution. This avoids possible synchronization effects, yielding to unbiased samples [10]. Despite having higher predictability, the uniform distribution is also used to bound the interval between samples, speeding up the convergence of the estimation resulting from sampling.

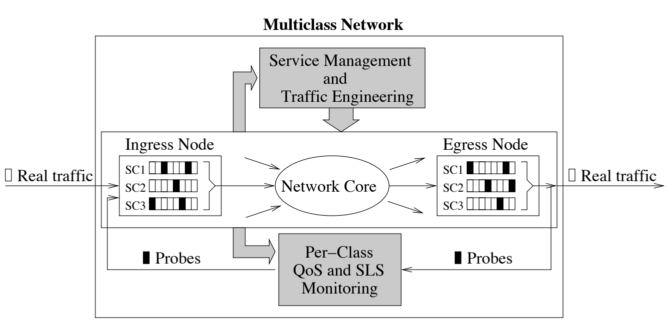

Figure 2: Edge-to-edge multiclass active monitoring

3.2 Enhancing QoS Monitoring

According to [7], commonly accepted and used probing patterns may fail to capture simultaneously common performance metrics in terms of shape and/or scale, when performing per-class active monitoring. This has motivated exploring alternative or complementary probing features capable of increasing multipurpose active monitoring efficiency.

In this context, taking into account the test results in Section 4, a versatile Back-to-Back On-Off probing source ($B2B_P$) with a deterministic On period and an Off period either deterministic or regulated by an Exponential, Pareto or Uniform distribution has been developed and tested. This source allows to regulate the number of probing samples per second and the time gaps between probing packets during each sample (see example in Figure 3). The resulting back-to-back probing streams aim at increasing probing sensitivity to queue variations by reducing the interpacket time between consecutive probes, while remaining simple and light.

Moreover, a new approach of coloring probes, i.e., exploring the effect of marking probes with different drop precedences, using single color or interleaved color schemes, has also been developed and explored (see also Figure 3). This approach aims at exploring AQM actions in case of queue congestion and the different probabilities of packets reaching the network boundary. This may be particularly useful to sense packet loss in the network domain.

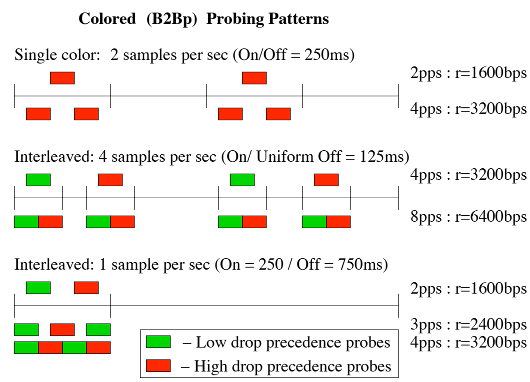

Figure 3: Proposed probing scheme

4. EVALUATING PROBING PATTERNS

The main objective of the tests is to assess the feasibility of multipurpose probing process in a multiclass network environment. Both the characteristics of probing patterns and the probing ability to capture each class behavior are studied aiming at a multipurpose QoS estimation. To pursuit this objective, several alternative probing schemes are explored varying the probing distribution, rate and drop precedence (color), in a per-class basis. In the analysis, for each service class, the probing measurement outcome is cross-checked against the corresponding measures using the real traffic within each class. This verification process is carried out resorting to a direct comparison of graphical results and statistical analysis of collected measurement samples.

4.1 Test platform

A simulation prototype of a multiclass network domain

Table 1: Service Classes SC_i

SC_i	Serv. Type	Prot.	$P_{i,p}$	$T_{i,p}$	Example	Traffic Src	i.a.t.	hold.t	Probing Characteristics
SC1	guaranteed (hard-RT)	UDP	IPTD ipdv IPLR	35ms 1ms 10^{-4}	VoIP Cir.Emulation Conv. UMTS	Exp. or Pareto on/off (64kbps, pkt=120B on/off = 0.96/1.69ms)	Exp. 0.3s	Exp. 90s	Distrib.: POI_p, $B2B_p$, FTP_p Mean Rate: 2 to 192 pps
SC2	predictive (soft-RT)	UDP	IPTD IPLR	50ms 10^{-3}	audio/video streaming	(256kbps, pkt=512B on/off = 500/500ms)	0.5s	120s	Color: Green, Red, Interleaved Packet Size: 100 bytes
SC3	best-effort	TCP	IPLR	10^{-1}	elastic apps.	FTP traffic (pkt=512B)	0.5s	180s	

was set up using NS-2. In this prototype, distributed edge-to-edge on-line monitoring is carried out at network edges to guide the dynamic admission of flows entering the multiclass domain [8]. In the test platform, three service classes (SCi) were defined and configured according to their QoS requirements (see Table 1). For each SCi, the defined thresholds $T_{i,p}$ for each QoS parameter $P_{i,p}$ (IP Transfer Delay (IPTD), IP Delay Variation (ipdv), IP Loss Ratio (IPLR)) consider the domain's characteristics and perceived QoS upper bounds for common applications and services.

The network domain consists of ingress routers I_1, I_2, a multiclass network core and an egress router E_1. The service classes SC1, SC2 and SC3 are implemented in all domain nodes, using a hybrid Priority Queuing - Weighted Round Robin (PQ-WRR(2,1)) scheduling discipline, with RIO-C as AQM mechanism. The domain internodal links capacity is 34Mbps, with a link delay of 15ms. The measurement time interval Δt_i is set to 5s.

4.2 Per-class measurement analysis

A per-class analysis of measurement outcome using Poisson probing (POI_p) with different rates (Test1) showed that:

(i) SC1: this is a very stable traffic class with IPTD and ipdv tightly controlled without suffering any loss, thus, probing is able to approximate this behavior for a probing rate as low as 2pps. In more detail, the estimated probing values are delimited within a range of 0.1ms around IPTD and ipdv class estimates. This is obviously a small variation in scale, in particular for IPTD. However, the shape of probing and class estimates is not well adjusted. As Figure 4 (Test1 @ 24pps) shows, increasing probing rate to 24pps clearly brings probing close to the class shape, with an ipdv overestimation upper bound of 0.05ms. A similar improvement in IPTD shape and scale is only met for higher probing rates. For instance, at 192pps, ipdv and IPTD probing estimations differ in 0.01ms from the class. Despite that, the intrusion cost introduced may be too high to be considered.

(ii) SC2: considering the less strict nature of this service class, for a probing rate of 2pps, IPTD is also fairly well captured (see Figure 4). However, a significant decrease on ipdv overestimation is only achieved for the higher probing rates under test (see example in Figure 5(b,b')). IPLR also benefits from probing rate increase as loss is gradually better detected, although a deficit on IPLR scale estimation of approximately one order of magnitude is still present.

(iii) SC3: for the lower probing rates, the behavior of SC3 is clearly the worst captured as regards the metrics under control. While IPTD and ipdv benefit from a probing rate increase, IPLR events are completely missed even at 192pps.

4.3 Improving multimetric estimation

Coloring probe is proposed and studied here as a method of improving IPLR estimation. Changing the drop prece-

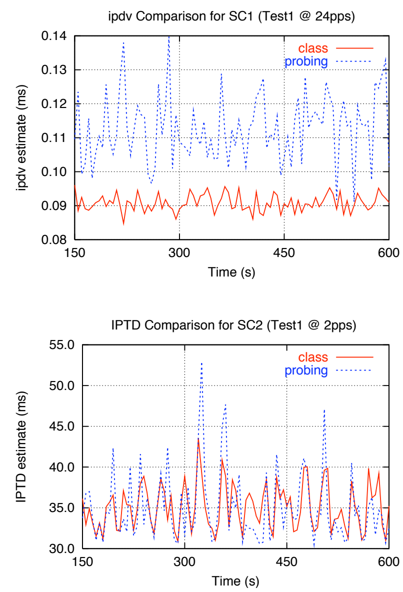

Figure 4: Comparison of class and probing measurements for: SC1 (24pps) and SC2 (2pps)

dence of probing packets from low to high (Test2), i.e., coloring probes from green to red, improves IPLR estimation significantly even for a probing rate of 2pps. This improvement is reflected in a better detection of class loss events, however, IPLR scale is clearly overestimated (see Figure 5 (a)) and IPTD slightly underestimated. This behavior is justified by AQM action on probes' precedence. In fact, when queue congestion increases, red probes are the first to be dropped and previously high-delayed green probes are now mostly discarded. Consequently, on average, IPTD estimation is slightly degraded, while loss detection is improved.

In order to improve the compromise between IPLR and IPTD estimation, probing patterns with an interleaved coloring scheme are tested, i.e., green and red packets are sent alternately (Test3). Comparing to Test1 outcome, a interleaved probing pattern of 2pps brings significant improvement on IPLR estimation for all traffic classes. Moreover, the degradation of IPTD estimation and IPLR overestimation noticed in Test2 is now much less pronounced. These positive results are further enhanced for higher probing rates where the accuracy of all metrics is increased. Generically, SC3 is the class that strongly benefits from interleaved probing, reaching a good compromise in the metrics' estimation for a probing rate of 24pps (see Figure 5(a')). Despite that, for SC2, IPLR still remains overestimated.

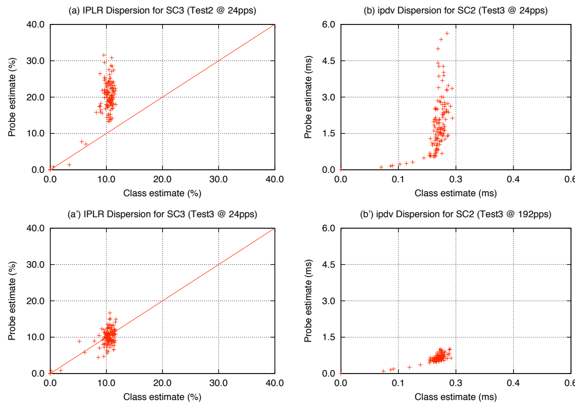

Figure 5: Dispersion of (a,a') IPLR for SC3 - red and interleaved @ 24pps; (b,b') ipdv for SC2 - interleaved @ 24 and 192pps

Regarding ipdv, as in Test1, ipdv estimates converge for probing rates above 96pps (see Figure 5(b,b')). Generically, it was found that ipdv is a rather sensitive metric to network load and its variability. The results illustrated in Test1 and Test3 show that ipdv scale for classes with high variability (SC2 and SC3) is difficult to obtain regardless of the test probing rate considered. While small ipdv variations are magnified by probing, ipdv mismatch under the different probing rates suggests that queuing delay oscillations persist across multiple time scales. For moderate loads, ipdv is more closely measured as the queues remain in a reasonable steady state. As decreasing the interpacket gaps through a probing rate increase did not successfully solve ipdv scale estimation, a $B2B_p$ probing pattern was considered. On average, this tries to keep regular and light the mean probing rate while reducing interpacket gaps.

Considering interleaved $B2B_P$ probing sources generating 8 and 24pps, respectively, with on and off periods of 125ms (corresponding to four back-to-back bursts), it was noticed that, when compared to POI_P with similar rates, interleaved $B2B_P$ probing streams lead to better estimates of ipdv scale, also with better results on IPTD and IPLR. This improvement is notorious in the experiment with probing at 24pps as shown in Figure 6, where the correlation coefficient between class and probing outcome is evaluated for all QoS metrics and all test scenarios previously considered.

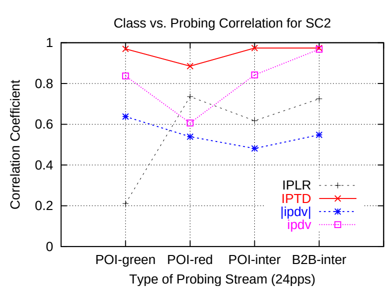

Figure 6: Correlation between class and probing estimates for different probing streams

5. CONCLUSIONS

In this paper, we have discussed the problematic of QoS monitoring in multiclass networks, pointing out directions to improve the simultaneous estimation of multiple QoS metrics. The evaluation tests have demonstrated that, although a probing rate increase generically improves QoS metrics' estimation accuracy, this improvement is rather dependent on each service class characteristics and on the type of metrics being estimated. In fact, even high probing rates were unable to match both the scale and shape of metrics such as IPLR, compromising at same time an acceptable trade-off between the estimation accuracy and probing overhead. In this context, a new hybrid and flexible on-off probing source, allowing to control the number of probing events and probing gaps, has been proposed. The resulting probing streams, especially conceived to improve ipdv estimation, when complemented by a proper coloring scheme can also improve the estimation of the other QoS metrics considered.

6. REFERENCES

[1] H. Asgari, P. Trimintzios, M. Irons, R. Egan, and G. Pavlou. Building Quality-of-Service Monitoring Systems for Traffic Engineering and Service Management. *Journal of Network and Systems Management*, 11(4):399–426, Dec. 2003.

[2] J. Coppens, et al. SCAMPI: A Scalable and Programmable Architecture for Monitoring Gigabit Networks. In *6th IFIP/IEEE MMNS'03*, Sept. 2003.

[3] J. Corral, G. Texier, and L. Toutain. End-to-end Active Measurement Architecture in IP Networks (SATURNE). In *PAM'03*, 2003.

[4] F. Georgatos, F. Gruber, D. Karrenberg, M. Santcroos, H. Uijterwaal, and R. Wilhelm. Providing Active Measurements as a Regular Service for ISPs. In *PAM'01*, Apr. 2001.

[5] K. Glossbrenner. Internet Protocol Data Comm. Service - IP Packet Transfer and Availability Performance Parameters. ITU-T I.380, 1999.

[6] S. Leinen and V. Reijs. Geant D9.7 - Testing of Traffic Measurement Tools. Geant Project, Sept. 2002.

[7] S. Lima, P. Carvalho, and V. Freitas. Tuning Active Monitoring in Multiservice IP Networks. In *Perf. Modelling and Evaluation of Heterogeneous Networks (HET-Nets'04)*, U.K., July 2004.

[8] S. R. Lima, P. Carvalho, and V. Freitas. Self-adaptive Distributed Management of QoS and SLSs in Multiservice Networks. In *IEEE/IFIP IM 2005*, Nice, France, May 2005. IEEE Press.

[9] B. Lowekamp. Combining Active and Passive Network Measurements to Build Scalable Monitoring Systems on the Grid. *ACM Performance Evaluation Review*, 30(4):19–26, 2003.

[10] V. Paxson, G. Almes, J. Mahdavi, and M. Mathis. Framework for IP Performance Metrics. RFC 2330 (Informational), May 1998.

[11] R. Prasad, C. Dovrolis, M.Murray, and K. Claffy. Bandwidth Estimation: Metrics, Measurement Techniques and Tools. *IEEE Network*, 2003.

[12] R. Whitner, G. Pollock, and C. Cook. On Active Measurements in QoS-Enabled IP Networks. In *PAM'02, Fort Collins CO*, Mar. 2002.