Practical design constraints for measuring utilization in hybrid paths using delay measurements

Jose Nuñez-Martínez†, Marc Portoles-Comeras†, Albert Cabellos-Aparicio‡
Josep Mangues-Bafalluy†, Jordi Domingo-Pascual‡
†Centre Tecnològic de Telecomunicacions de Catalunya (CTTC) Castelldefels (Barcelona), Spain
{jnunez, mportoles, jmangues}@cttc.cat
‡Universitat Politècnica de Catalunya Barcelona, Spain {acabello,jordid}@ac.upc.edu

***Abstract*— Recent research results have shown how measurements of bandwidth metrics using traditional tools and techniques have to be reconsidered in the presence of WLAN links. The main reason behind this is the CSMA/CA protocol used to regulate the distributed access in the wireless medium. In this paper we identify a set of practical issues that tools based on delay measurements must face to implement end-to-end bandwidth metrics in hybrid paths (i.e., wired/wireless). The paper provides results through extensive simulation to validate the findings of the study.**

***Index Terms*—path utilization, wired/wireless, bandwidth metrics, delay measurements**

I. Introduction

The CSMA/CA protocol used to regulate channel access in WLAN networks invalidates some of the traditional assumptions used to design bandwidth measurement tools [1]. Recent research studies, such as the ones presented in [3] and [4] have revisited the use of delay-dispersion measurements to infer bandwidth metrics in the presence of CSMA/CA links. The results mainly show that such techniques target the *Achievable Throughput* (AT) metric, rather than the *Available Bandwidth* (AB) and that the transient stage that probing packets suffer impact the measurements by biasing the results.

This paper takes as a reference the Forecaster tool [2] and reviews the use of delay samples to measure the utilization of end-to-end hybrid path (wired/wireless) scenarios. To this end, the paper revisits the use of delay samples in the presence of CSMA/CA regulated links, reveals some steady state properties of this type of measurements and discusses practical design constraints related to their use.

The contributions of the paper are multiple-fold. First, it reveals and provides arguments to sustain that delay-based measurements also target the Achievable Throughput (AT) in the presence of CSMA/CA links. Second, it shows that delay-based measurements are also affected by the transient regime characterizing the interaction of flows in the WLAN link. Third, the paper reveals that the use of delay thresholds to differentiate delay samples is a desirable feature to build accurate tools that can operate with a reasonable amount of probing samples. Finally, the paper also shows how linear projections, such as the one proposed in Forecaster, do not hold anymore in the presence of CSMA/CA links and alternative iterative methods are advisable.

The remainder of this paper is structured as follows. Section II describes the principles behind delay-based measurements of utilization based on using Poisson sequences. Section III studies the implications of using such techniques in the presence of WLAN links (with CSMA/CA access protocol). Section IV studies which are the implementation issues related to building delay-based measurement tools in end-to-end paths with WLAN links. Section V provides validation of the paper statements through simulation results of practical cases. Finally, section VI presents related work and section VII concludes the paper.

II. Model of the inter-relation between delay and cross-traffic

This section introduces a model of the inter-relation existing between delay measurements and the utilization of a network link from a FIFO queuing perspective.

A. The utilization of a network link: single-hop FIFO case

Let us define the *hop-workload process* $\{W(t), t \geq 0\}$ (see [6]) as the sum of the service time of all packets present in a FIFO queue and the remaining service time of any packet that may be in service at time *t.*

Assuming that packets are served in a work-conserving fashion, we can define the process $U(t)$ describing the utilization of the network link associated to this FIFO queue as,

$$U(t) = \begin{cases} 1 & W(t) > 0 \\ 0 & W(t) = 0 \end{cases} \tag{1}$$

Then, the long term average utilization of the link can be expressed as

$$u_l = \lim_{t \to \infty} P[U(t) = 1]. \tag{2}$$

B. The packet probing process

Let us now define a probing sequence, consisting in a series of packets of length L that are sent at instants $\{T_n\}$ through the targeted network link.

Under the terms of this section, we define the probing rate (r_p) as follows,

$$r_p = \frac{L}{E[T_n - T_{n-1}]}. \tag{3}$$

The rest of the study assumes that the probing sequence follows a Poisson distribution with parameter λ. That is,

$$P[T_n - T_{n-1} < t] = 1 - e^{-\lambda t}.$$

C. Single-hop utilization: Aggregation of probing and cross-traffic.

The utilization of a network link, when we introduce a probing flow, is the one that results from the aggregation of this probing flow and the cross-traffic that traverses this same link.

Let us denote as u_c the long term average cross-traffic utilization of a network link. Let us use the term u_p to denote the probing flow utilization of this same network link when probing at rate r_p. Then,

$$u_l = \lim_{t\to\infty} P[U(t) = 1] = u_c + u_p. \quad (4)$$

D. Delay measurements and the sampling process

We construct now a sampling process of the utilization of a network link $U(t)$ using the observation that delay measurements can be used as indicators of such utilization.

Taking the probing sequence into account, we denote as $\{d_n\}$ the sequence of delays that each one of the probing packets suffer when they traverse the network link and $d_{min}$the minimum delay they can experience. Using this, we can define a sampling process of the utilization such that,

$$U_n = U(T_n) = \begin{cases} 1 & d_n > d_{min} \\ 0 & d_n = d_{min} \end{cases} \quad (5)$$

The *Poisson* distribution used to send the probing flow guarantees that,

$$\lim_{n\to\infty} P[U_n = 1] = \lim_{t\to\infty} P[U(t) = 1] = u_c + u_p \quad (6)$$

This observation leads to stating that delay samples can be used to infer the utilization of a network link. Specifically, the probability $P[d_n > d_{min}]$ constitutes an indicator of the utilization of a network link.

E. Measuring bandwidth metrics using delay samples

In wired scenarios, the authors of [2] show how, when using *Poisson* probing and assuming FIFO queuing along the whole path, finding the probing rate r_p that causes expression (6) to reach 1 is equivalent to finding the *Available Bandwidth* (*AB*) of the network path. Formally, it may be expressed as

$$AB = \lim_{n\to\infty} inf\{r_p : P[d_n > d_{min}] = 1\}. \quad (7)$$

Furthermore, this study also shows that when probing at rates below *AB*, there exists a linear relationship between the value of $P[d_n > d_{min}]$ and the probing rate r_p. Interestingly, the authors show how the linear approximation still holds in multihop scenarios.

Based on this mode the authors proposed a tool named *Forecaster* that exploits this linear relationship to infer bandwidth metrics, and that minimizes traffic intrusiveness with respect with traditional tools (e.g., pathload).

III. THE RELATION BETWEEN BANDWIDTH METRICS AND DELAY MEASUREMENTS IN CSMA/CA LINKS

This section introduces the issues faced when using the model described in Section II in the presence of wireless links.

A. Reconsidering the relation between utilization and delay samples in the presence of CSMA/CA links

Practical CSMA/CA scheduling, such as the one implemented in IEEE 802.11 devices, breaks the work-conserving assumption used in (1). When applying the DCF mechanism, stations may have packets waiting to be transmitted (i.e., waiting for the backoff to expire), even though the medium is idle.

From a delay sampling perspective this observation has two implications on the measurements. On one side, while waiting for a backoff to expire, the probing queue may hold packets even when the medium is not being used. Delay samples obtained with such packets might produce erroneous indications that the system is being used.

On the other side, a probing station may obtain d_{min} samples even when contending stations have packets waiting to be transmitted. From a measurement perspective, such delay samples are correctly measuring the idle/busy state of the channel. However, this has implications in the ability to infer certain bandwidth metrics when using delay samples.

An important note related to the above statement is that the probability $P[d_n > d_{min}]$ can only reach 1 when all probing packets are queued before transmission. Otherwise, there is a non-zero probability that they are transmitted before any other packet in the system, even when contending stations are in backlog.

B. Bandwidth metrics in presence of CSMA/CA links using delay samples

The authors of [3] and [4] have recently shown how in the presence of CSMA/CA regulated channel access, *delay-dispersion* based bandwidth measurement techniques do not target the AB but rather the *Achievable Throughput* (AT)[1]. Under the conditions in this paper, AT is defined as

$$AT = \lim_{n\to\infty} sup\left\{r_p : \frac{r_o}{r_p} = 1\right\}, \quad (8)$$

where r_o is the rate at which the probing flow is received once it has traversed the measured link. The studies presented in [3] and [4] show how, once achieving steady state conditions, dispersion measurements are suitable indicators of rate responses and can be used to infer (8).

Figure 1 illustrates how this same statement also holds in the case of using delay samples rather than delay-dispersion samples. The Figure plots the utilization measured (i.e., $P[d_n > d_{min}]$), when probing an 802.11 link at different probing rates (r_p in the x axis). The Figure also plots the evolution of the throughput that a contending station observes when trying to inject a 4Mbps flow into the wireless link. To obtain this curve, a probing station is configured to transmit at a PHY rate of 11Mbps. There is a single contending station (also configured at 11Mbps) that is transmitting packets of

[1] The authors of [4] call this metric *fair-share*, which refers to the share of channel access that the probing node can get. However, reference [3] generalizes the same metric to a FIFO+CSMA/CA access model.

1500 bytes at 4Mbps. The Figure is obtained using probing sequences of more than 20k packets[2] . Please note that all the experiments depicted in this paper have been also obtained using NS-3 [10].

Figure 1 shows that the measured utilization only reaches 1 when the probing rate corresponds to the AT. Moreover, when the contending flow starts being affected by the probing flow (i.e., when AB is reached), the measured utilization is still below 1.

The reason behind this observation lies in the specific implementation of the CSMA/CA mechanism associated to the IEEE 802.11 spec. As explained before, the non work-conserving property of the DCF mechanism causes that $P[d_n > d_{min}]$ strictly equals 1 only when the probing flow suffers queuing (with probability 1) before being transmitted. This only happens when the probing station enters backlog, i.e., when it transmits above its AT.

As a consequence, the observation in Figure 1 can be generalized to any IEEE 802.11 single hop scenario. We can state that the AB is not the metric targeted using delay-based measurements of the utilization but rather the AT. Then, when using delay samples:

$$AT = \lim_{n\to\infty} inf\{r_p : P[d_n > d_{min}] = 1\}. \quad (9)$$

Finally, Figure 1 also illustrates that the relationship between the measured utilization (i.e., $P[d_n > d_{min}]$) and the probing rate (when probing below the AT) is not linear anymore (as in the FIFO case). This observation prevents the use of techniques such as the linear projection used in *Forecaster* to measure bandwidth metrics in the presence of WLAN links. As a consequence, iterative probing techniques are advisable to find the AT.

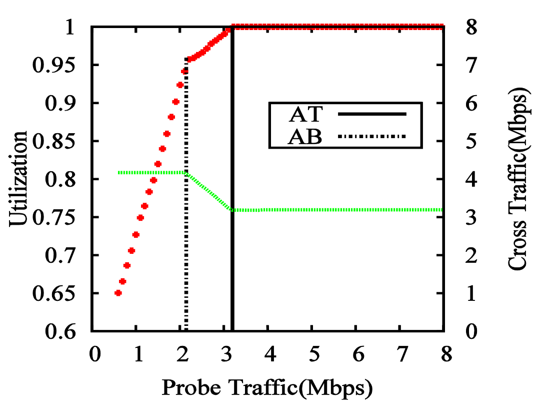

Figure 1 Utilization measurements versus the probing rate. The measured utilization reaches 1 when the probing reaches the AT.

IV. PRACTICAL ESTIMATION OF THE UTILIZATION IN THE PRESENCE OF WLAN LINKS

This section discusses practical considerations to take into account to build a delay-based bandwidth measurement tool to be used in wired/wireless scenarios. The section discusses the implications of having probing sequences with a limited number of samples and alternatives to increase the accuracy of the measurements.

A. *Estimation using a limited number of samples*

The previous section described the inter-relation between the use of delay samples to measure utilization and bandwidth metrics in the presence of CSMA/CA links. Following the idea of *Forecaster,* we study the possibility to use the following estimator of the utilization in the presence of CSMA/CA regulated links:

$$\hat{u} = \frac{1}{N}\sum_{n=0}^{N-1} 1_{\{d_n > d_{min}\}}. \quad (10)$$

A key parameter of equation (10) is the number of samples N used to obtain the estimate. Figure 2 illustrates, for the same scenario as in Figure 1, which is the relation between the expected measurement of the AT using (10) and the number of samples used. In order to produce the estimates of the Figure, we replicate the simulation more than 100 times and use as an estimate of the AT, the first probing rate that produces a measure of 1.

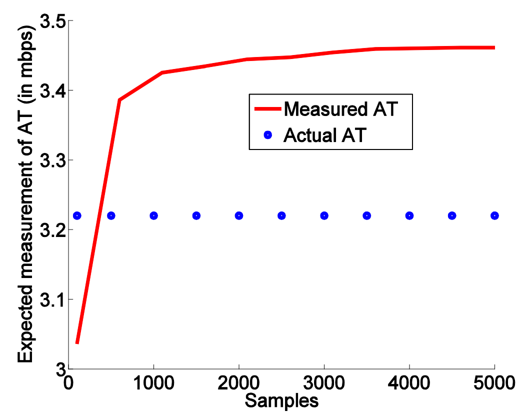

Figure 2 Expected measured AT using (10) with a variable number of delay samples (*N*)

Figure 2 illustrates two important features related to the practical use of (10). On one side, when probing with a low number of samples, the measure tends to underestimate the actual value of the AT. This is a common property of tools based on counting minimum-delay samples. Indeed, as the value of the utilization approaches 1 the probability of having samples of d_{min} reduces considerably (e.g., see Figure 1). Even more, the non-linear relationship between utilization estimates and probing rate intensifies this effect and may lead to high utilization measurements even when using a probing rate still far from the AT.

On the other side, as the number of samples increases, so does the expected measurement value. However, as it can be seen in Figure 2, when the number of samples is sufficiently large, the AT is overestimated. The origin of this effect lies in the transient state behavior that delay samples experiment before reaching steady state conditions [3].

The following subsections describe two design options to be

[2]The reason behind using such a large number of samples is that we are interested in the system response once in steady state conditions. In fact the measure is obtained after removing the first 2000 delay samples to ensure no distortion from any transient stage.

used for dealing with the two issues illustrated by Figure 2.

B. Removal of the first samples from the measurement

This section deals with the overestimation caused by the presence of the initial transient state of delay samples. In a previous work [3], we studied this effect in relation to CSMA/CA systems. The results showed how the first delay samples of a probing train constitute biased samples of the actual distribution of delays once in steady state.

From a practical perspective, removing these first samples from the measurement may improve the accuracy of the measurement process. Expression (10) can be modified to account for this as follows:

$$\widehat{u_\delta} = \frac{1}{N-\delta}\sum_{n=\delta}^{N-1} 1_{\{d_n > d_{min}\}}. \qquad (11)$$

Figure 3 illustrates this. It shows how, by removing the first samples, the AT estimation approaches the actual measure. However, the Figure also illustrates how, in certain cases, the amount of samples to remove in order to obtain accurate measurements can be rather large.

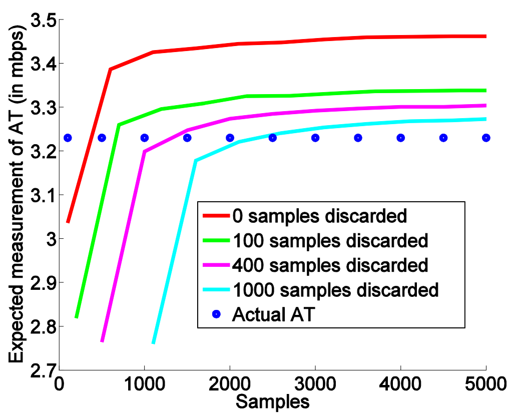

Figure 3 The effect of removing the first delay samples from the estimation of the utilization.

However, as the Figure 3 shows, the relation between the number of samples and AT measurements presents similar slope, regardless of the number of initial samples removed. The consequence of this is that there exists a trade-off between the number of probing packets to use (intrusiveness) and the accuracy that can be expected from the measurement.

C. Using a delay threshold rather than an absolute minimum delay value

This section analyzes the possibility to use delay thresholds rather than absolute minimum delays to classify delay samples. This is a common strategy to use in practice when there is no fine grained synchronization between end hosts, for example when using drift removal tools to obtain relative synchronization [7].

Let us propose the use of the following modification to the estimator of the utilization of a path,

$$\widehat{u_\tau} = \frac{1}{N-\delta}\sum_{n=\delta}^{N-1} 1_{\{d_n > d_{min}+\tau\}}, \qquad (12)$$

where τ is a delay threshold added to the minimum-delay value. The estimator as presented here, counts as samples of d_{min} all those delay samples that are lower than the threshold.

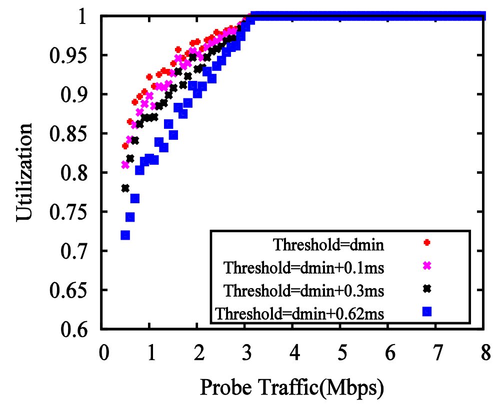

Figure 4 Utilization curve in case the bottleneck is in the wireless part of the end-to-end path. Utilization reaches 1 at the same probing rate (AT) irrespective of the delay threshold used.

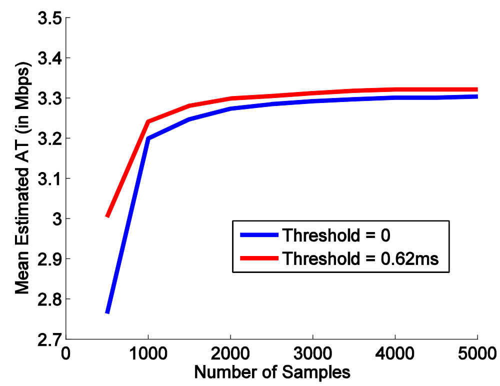

Figure 5 Comparison of the estimated AT when using two different delay thresholds and a varying number of delay samples. To obtain the Figure we remove the first 400 samples of the measure

Figure 4 shows, the effect of computing the utilization curve (i.e., measurements of $P[d_n > d_{min}]$) using (12) with different values of the delay threshold. As it can be seen, the use of a delay threshold does not affect the possibility to identify the AT. As the Figure reveals, the measured utilization reaches 1 when using the probing at the AT, regardless of the threshold employed.

Additionally, the Figure illustrates one important advantage of using the delay threshold. Delay thresholds sharpen the slope of the utilization curve. As a consequence, the probability of obtaining a false positive of the AT is lower.

However, the use of delay thresholds presents two specific issues that limit the advantages of their applicability, illustrated by Figure 5. The plot compares the estimated AT (following the same strategy as in Figure 2) when using no delay threshold and a threshold of $620\mu s$.

First, the threshold produces an overestimation of the estimated AT. This overestimation can be made arbitrarily low when an appropriate number of initial samples are removed from the probing sequence.

Second, Figure 5 reveals that the amount of samples required to produce accurate measurements (i.e., within a limited bound of the actual value) is approximately the same regardless of the threshold used. An alternative lecture of the Figure is that, in absolute terms, the maximum expected error of a measurement can be bounded using delay thresholds, irrespective of the number of samples used to obtain it.

D. End to end measurement of utilization in hybrid wired/wireless

As a summary of the observations in sections III and IV there are a number of issues to take into account in order to develop end-to-end bandwidth measurement tools that are resilient to the presence of CSMA/CA (WLAN) links.

First, the metric targeted through delay measurements is the Achievable Throughput (AT). Even in the presence of FIFO (wired) links, the utilization only reaches 1 when probing at the AT (with a sufficiently large number of samples).

Second, the non-linear behavior of delay-based measurements of the utilization prevents the use of projection mechanisms such as the one used in *Forecaster*. Instead iterative mechanisms based on algorithms such as the binary search (used in pathload) are advisable. Third, we have provided arguments to show how using a threshold to compute utilization (see expression (12)), can help improving the accuracy of the measurement when using a limited number of samples or inaccurate measurements of delay. As shown below in the validation section, the use of this threshold does not impact the accuracy of measurements when wired (FIFO) links are present. However, as shown also, as the number of delay samples increases the benefits of using delay thresholds rapidly diminish.

Finally, as shown above, enlarging the number of delay samples used does not help reducing the overestimation caused by initial transient delays. Instead, as Figure 3 reveals, it is recommended to remove the first samples when obtaining estimates of the utilization. However, as the Figure also shows, the amount of samples required to obtain accurate measurements of bandwidth can be rather large.

V. VALIDATION

This section validates our findings through NS-3 [10] simulation. In short, we aim to show that it is feasible to infer bandwidth metrics in end-to-end wired/wireless paths under the set of constraints described before.

A. Scenario setup

Figure 6 depicts the simulated wired/wireless path used to validate our findings. The hybrid path includes 4 FIFO wired links and one wireless link. The wired links are configured at 10Mbps (node 1 to node2), 100Mbps (node 2 to node 3), 1Gbps (node3 to node4) and 10Gbps (node 4 to AP) respectively. The last hop of the hybrid path (AP) is an 802.11 WLAN link with multiple WLAN nodes transmitting cross-traffic at different rates.

In all the simulated cases, node 1 transmits sequences of *Poisson* probing traffic towards node 5. Each wired link is loaded with UDP flows representing the cross-traffic. Regarding the wireless link, we have considered two different scenarios. A simple scenario where only one contending node transmits cross-traffic (at 5Mbps) towards de AP, and a more complex one with 3 contending nodes sending traffic at 250kbps, 500kbps and 2Mbps respectively. Moreover each node uses a different packet size. Specifically, packets size of 200, 500 and 1500 bytes were chosen.

B. Simulation results

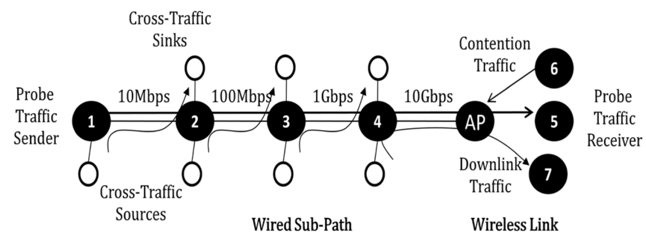

Figure 6 Hybrid scenario under evaluation

First, we focus on the simple case, in this scenario the bottleneck of the path is the wireless link. In this case, the contending node transmits cross-traffic at 5Mbps.

Figure 7 plots the measured utilization vs. the probing rate. Precisely, the utilization reaches 1 at ~3.2Mbps. This point represents the fair-share (AT) of the channel bandwidth corresponding to the probing flow. Also, the plot shows that the estimation is irrespective of the threshold used to classify delay samples. Even more, the Figure shows that using a larger threshold the knee is sharper, and this may allow to estimate the AT using a limited amount of probing packets.

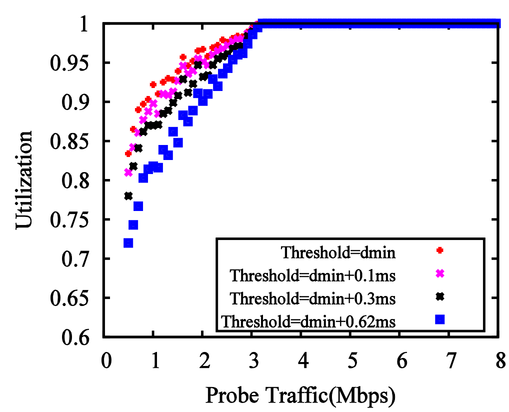

Figure 7 Utilization curve in case the bottleneck is in the wireless part of the end-to-end path. Utilization reaches 1 at the same probing rate (AT) irrespective of the delay threshold used.

Figure 8 plots the results obtained for the simple case, in this case the bottleneck is at the wired part of the path. In particular, the first link is overload and has 1.5Mbps of spare capacity. The Figure shows that the AT of the path (the AB in this case, since it is a FIFO link) is accurately estimated, even when using delay thresholds (up to 620µs).

Finally, Figure 9 plots the results for the complex scenario. In this case the bottleneck is again in the WLAN link, which includes 3 contending stations. The Figure shows that the estimation is accurate, and that again the thresholds improve the sharpness of the utilization curve. Again, this suggests that a tool may reduce the amount of samples used to estimate and hence, limit its intrusiveness.

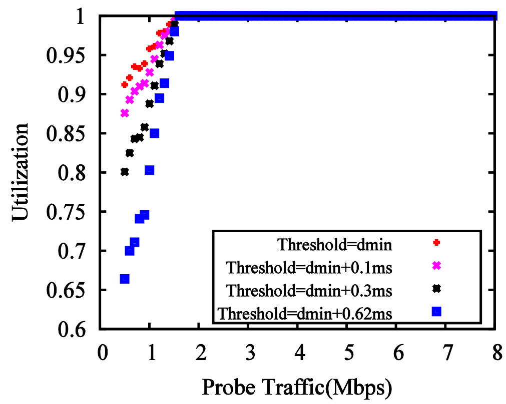

Figure 8 Utilization curve in case the bottleneck is in a wired (FIFO) link of the path. Utilization reaches 1 at the same probing rate (AT) irrespective of the delay threshold used.

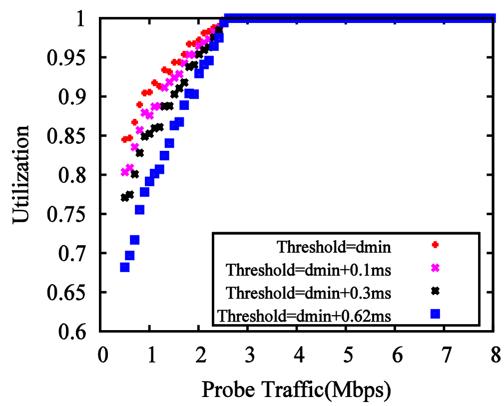

Figure 9 Utilization curve in a scenario with three WLAN contending nodes transmitting at different rates. The bottleneck is in the WLAN link (AT~2.5Mbps).

VI. Related Work

Beyond the classical bandwidth measurement literature based on the FIFO assumption and recent advances related to dispersion-based measurements in WLAN environments, there are a number of relevant approaches to the study presented in this paper.

First, the authors of [5] introduce *ProbeGap*, the first tool designed to infer bandwidth in a single wireless link. The tool is based on a heuristic that tries to find a "knee" in the cumulative distribution function of delays measured when transmitting a probing flow through a wireless link. In fact, this heuristic is a precursor of the estimator described by expression (12) as the authors try to find a delay threshold to classify delay samples and, then, estimate the utilization of the link. However, as the authors show, this heuristic fails when the values of the measured utilization are high.

Also, several authors have designed passive tools that "observe" the medium to account for the ratio of idle/busy periods ([8], [9]). These approaches are accurate to find bandwidth metrics of WLAN links. However they cannot operate on an end-to-end path basis and require, in some cases, access to low-layer parameters of the wireless devices.

VII. Conclusions

This paper revisits the use of delay measurements in order to infer end-to-end bandwidth metrics in hybrid paths (wired/wireless). It also presents a series of design considerations to take into account when developing bandwidth measurement tools based on delays samples for such scenarios.

The study reveals how delay-based measurements target the Achievable Throughput rather than the Available Bandwidth, just as reported previously for dispersion based measurements [3]. The paper also reveals how the transient behavior of the delays in CSMA/CA also leads to obtaining inaccurate measurements of bandwidth in this scenario, particularly with short probing trains.

Additionally, the study reveals how delay-based measurements are prone to considerable errors when the number of samples used in the estimations is limited. As a solution to this we have proposed the use of delay thresholds to classify delay samples into busy and idle. This relaxes the probability to have erroneous measures using short probing sequences without affecting the accuracy of the estimate.

Acknowledgements

This work was supported in part by the Ministry of Science and Innovation under grant number TEC2008-06826/TEC (project ARTICO). Additionally, it was also partially funded by grant PTQ08-01006439.

References

[1] M. Jain, C. Dovrolis, *"End-to-End Available Bandwidth: Measurement methodology, Dynamics, and Relation with TCP Throughput"'* in Proc. of ACM SIGCOMM, 2002.

[2] M. Neiginhal, K. Harfoush, H Perros, "*Measuring Bandwidth Signatures of Network Paths*", in Proc. of IFIP NETWORKING, 2007.

[3] M. Portoles-Comeras, A. Cabellos-Aparicio, A. Banchs, J. Mangues-Bafalluy, J. Domingo-Pascual, *"Impact of transient CSMA/CA access delays on Active Bandwidth Measurements",* in Proc. of ACM IMC, 2009

[4] M. Bredel and M. Fidler,*"A Measurement Study of Bandwidth Estimation in IEEE 802.11g Wireless LANs using the DCF"*, in Proc. of IFIP NETWORKING, 2008

[5] K. Lakshminarayanan, V. N. Padmanabhan, J. Padhye. *"Bandwidth Estimation in Broadband Access Networks"*, in Proc. of ACM IMC , 2004

[6] X. Liu, K. Ravindran, D. Loguinov, *"A Queuing-Theoretic Foundation of Available Bandwidth Estimation: Single-Hop Analysis"*, in IEEE/ACM Transactions on Networking, Vol. 15, No. 6, 2007

[7] D. Veitch, S. Babu, A Pasztor, *"Robust synchronization of software clocks across the Internet"*, in Proc. of ACM IMC, 2004

[8] M. Ahmad, Y. Khan, D. Veitch, *"Speedo: Realistic achievable bandwidth in 802.11 through passive monitoring"*, in Proc. of IEEE LCN, 2008

[9] S. H. Shah, K. Chen, and K. Nahrstedt, *"Available bandwidth estimation in ieee 802.11-based wireless networks"*, in Proc. of ISMA/CAIDA Workshop on Bandwidth Estimation, 2003

[10] The NS-3 network simulator, http://www.nsnam.org/