An Analysis of Active End-to-end Bandwidth Measurements in Wireless Networks

Andreas Johnsson, Mats Björkman and Bob Melander
{andreas.johnsson, mats.bjorkman, bob.melander}@mdh.se

The Department of Computer Science and Electronics
Mälardalen University, Sweden

***Abstract*— For active, probing-based bandwidth measurements performed on top of the unifying IP layer, it may seem reasonable to expect the measurement problem in wireless networks to be no different than the one in wired networks. However, in networks with 802.11 wireless bottleneck links we show that this is not the case.**

The results from the experiments presented in this paper show that the measured available bandwidth is dependent on the probe packet size (contrary to what is observed in wired networks). Another equally important finding is that the measured link capacity, using the well known TOPP model, is dependent on the probe packet size *and* on the cross-traffic intensity.

The underlying reasons for the observed differences are analyzed by incorporating the characteristics of 802.11 wireless networks into the TOPP model. The extended model is applicable to other end-to-end bandwidth measurement methods as well, such as BART, Pathload and PTR.

I. Introduction

Wireless networks, used when connecting to the Internet or when several nodes want to communicate in an ad-hoc manner, are becoming more and more popular. Because of the increased dependence on wireless network technology, it is important to ensure that methods and tools for network performance measurement also perform well in wireless environments. In this paper, we focus on active performance measurements in terms of network bandwidth, both link capacity and the unused portion thereof; the available bandwidth.

Measurement of network properties such as available bandwidth in a variety of networks are important for network error diagnosis and performance tuning but also as a part of the adaptive machinery of network applications such as streaming audio and video.

State-of-the-art active end-to-end bandwidth measurement methods are for example BART [1], Pathchirp [2], Pathload [3], PTR [4], Spruce [5] and TOPP [6][7]. The basic principle is to inject a set of measurement packets, so called *probe packets*, into the network. The probe packets traverse the network path to a receiver node, which time stamps each incoming probe packet. By analyzing these time stamps, estimates of the link capacity (a physical property of a link) and/or the available bandwidth (the unused portion of the link capacity) can be made. For many end-to-end available bandwidth measurement methods no prior knowledge of the underlying network topology is needed. That is, bandwidth estimation methods are well suited for end-to-end performance measurements and monitoring in all sorts of networks. The existing methods differ in how probe packet are sent (the flight patterns) and in the estimation algorithms used. A good overview of available bandwidth measurement methods and tools can be found in [8].

In this paper we study how the properties of wireless 802.11 bottleneck links affect estimates reported by active end-to-end bandwidth measurement methods. We measure, analyze and describe how the characteristics of bandwidth estimation change in wireless scenarios. By conducting experiments where we vary the probe-packet size we show that the obtained bandwidth estimates reported vary. Further, by varying the cross-traffic rate we illustrate how the estimated link capacity obtained by using the well known TOPP model is not constant.

To understand the obtained bandwidth estimates reported we incorporate the properties of 802.11 wireless networks into the TOPP model. Using the extended model we describe why and how the probe-packet size affect the bandwidth estimate. Further, we discuss the relation between the estimated link capacity (using TOPP) and the real link capacity of a wireless 802.11 bottleneck link.

The extended TOPP model is also applicable to, and will thus explain why estimates produced by other end-to-end bandwidth measurement methods (e.g. BART, Pathload and PTR) vary with the probe-packet size in wireless 802.11 networks. We discuss this in detail in the paper.

The bandwidth measurements have been performed in a testbed containing both wireless and wired hops. Our testbed topology only consist of one wireless 802.11 bottleneck link, to be used as an access link, since that is a common way to deploy wireless networks. To produce measurement results we have used DietTopp, a tool that implements the TOPP model, which measures the available bandwidth and link capacity of an end-to-end path. For comparisons and to illustrate that our observations concerning the available bandwidth estimates are not tied to a certain measurement tool, we have also used the tool Pathload, which also measure the available bandwidth of an end-to-end path, in our experiments.

Earlier work has touched upon the problem of end-to-end measurement of bandwidth in wireless networks. In [9] we discuss the main problem areas when deploying existing band-

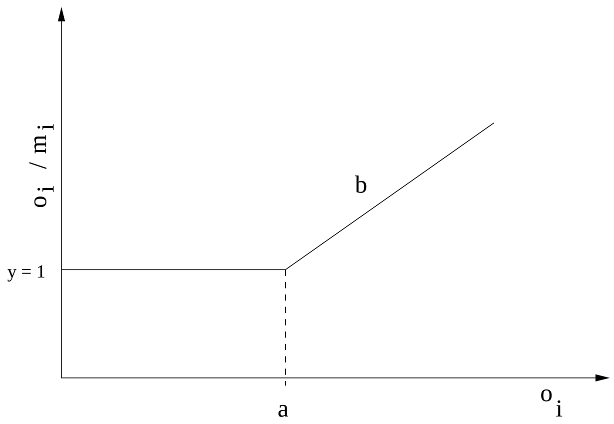

Fig. 1. Plot of the ratio o_i/m_i as a function of o_i.

width measurement methods in wireless networks. Further, measurement results presented in [10] indicate that the available bandwidth estimates is dependent of the probe packet size in multi-access networks, such as wireless 802.11 networks. The authors are not restricted to bandwidth measurements in 802.11 networks, but rather investigate a range of access network techniques.

This paper complement and extend the work made in [10]. In this paper we develop an analytical model describing the measurement results. Further, compared to [10] we use more complex testbed scenarios along with a different set of bandwidth measurement methods.

The rest of this paper is organized as follows. Section II-A describes the measurement model. DietTopp, which is our implementation of a simplified TOPP model, is also presented. Section II-B is a description of the testbed we have used for the investigation of the bandwidth measurement problem in wireless networks. In Section II-C the experiments are described along with the research question. Section III show measurement results from using DietTopp in wired as well as in wireless networks. We discuss the results and compare them to results obtained by Pathload. In Section IV we extend the original TOPP model to describe and explain the obtained measurement results. Further, in Section V some important observations are made. The paper ends with conclusions in Section VI.

II. Experimental setup

This section describes the experimental setup used in order to study the behavior of bandwidth measurement methods in wireless 802.11 networks. That is, the measurement model and the tools used in the experiments. The testbed setup and finally what kind of bandwidth measurements that have been performed along with their relevance are also presented.

A. Measurement model

The TOPP model [6] is well suited to describe and explain how end-to-end bandwidth measurement methods function [11]. Therefore we use that model to study the impact of

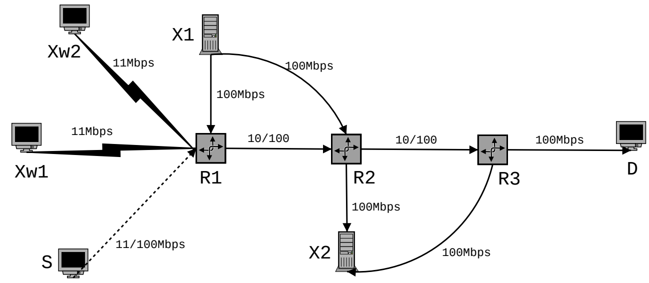

Fig. 2. The testbed is constructed by one wireless link, three routers and several cross-traffic generators (on both the wireless and the wired links)

wireless bottleneck links on bandwidth estimates. The TOPP model is described below.

The available bandwidth on a single link can be described as the unused portion of the link capacity during some time period. The end-to-end available bandwidth is the minimum available bandwidth for each link in the path. That link defines the bottleneck.

The basic assumption is that the probe packets gain a proportional share of the link capacity when traversing the bottleneck link. That is

$$m = \begin{cases} m & if m < a \\ \frac{o}{(x+o)} * L & if m \geq a \end{cases} \quad (1)$$

where m is the measured probe-packet rate, o the offered probe-packet rate, x the cross-traffic rate on the link with capacity L and a is the available bandwidth on the bottleneck link.

Equation 1 can be converted into a linear form describing the quotient o/m

$$\frac{o}{m} = \begin{cases} 1 & if m < a \\ (1 - \frac{a}{L}) + \frac{o}{L} & if m \geq a \end{cases} \quad (2)$$

where a is the available bandwidth on the specific link. A common definition of a is $a = (L - X)/\delta$ where X is the amount of cross traffic that flows through the link during time δ. If plotting Equation 2 in a diagram the result will be much like the theoretical one in Figure 1. This theoretical curve is in [11] called the *rate response curve*, that is, the relation between the offered rate and the measured rate. We adopt that term in this paper. Further, in [11] it is shown that the fluid-like analysis described above can give erroneous bandwidth estimates when taking packet-level interactions in the router queues into account. Especially if the cross traffic is bursty (e.g. Pareto distributed) or if there exist several secondary bottlenecks [6]. By using longer probe-packet trains instead of probe-packet pairs the obtained rate response curve asymptotically moves towards the fluid curve shown in Figure 1. The tools used in the experiments use long probe-packet trains. Further, we study the effect of one wireless bottleneck link on the rate response curve. Thus, for the objectives in this paper we believe that the original fluid model is sufficient.

To study how wireless 802.11 networks affect available bandwidth estimates (i.e. how the rate response curve in Figure

1 is affected) we have implemented a tool called DietTopp [12]. (In [7] preliminary results obtained from DietTopp in wireless networks is discussed along with the tool itself.) It measures the end-to-end available bandwidth along with the link capacity of the bottleneck link. Further, Pathload has been used to show that the results are not tied to a certain tool.

Both DietTopp and Pathload directly or indirectly use the information embedded in rate response curve to estimate the available bandwidth. We describe how below:

DietTopp injects a set of probe-packet trains at an increasing rate in the interval $[o_{min}, o_{max}]$. On the receiver side each probe packet is time stamped in order to calculate m_i for each incoming probe-packet train i. The probe-packet train rate increases for each successive train, hence the bottleneck link will be congested at some point (corresponding to $o_i = a$ in Figure 1). When all probe-packet trains have traversed the network path the quotient o_i/m_i can be plotted. The rate response curve in Figure 1 is used as an example. DietTopp uses linear regression to estimate the linear segment b. The end-to-end available bandwidth is defined as the offered rate corresponding to the intersection of b and $y = 1$. Further, the slope of b corresponds to the bottleneck link capacity according to [6].

To speed up the probing phase of DietTopp it is desired to avoid measurements with an offered rate o below a. That is, DietTopp wants to ensure that $o_{min} > a$. This is done by estimating m_{max} which is done by injecting a set of probe packets at rate o_{max} (could be the link capacity of the access link for example) and then measure their separation at the receiver. According to [6] m_{max} is greater than the available bandwidth (m_{max} is referred to as the asymptotic dispersion rate in [13]). DietTopp also assumes only one bottleneck link between the end nodes contrary to the more computational expensive TOPP model.

Pathload is based on observations of the one-way delay of probe-packet trains. If the offered probe-packet rate is above the available bandwidth the one-way delay will show an increasing trend. On the other hand, if the one-way delay doesn't show an increasing trend the offered rate is below the available bandwidth. To locate the available bandwidth, Pathload deploys binary search. That is by varying the offered probe rate and investigating the one-way delay. In [11] it is stated that there is a strong statistical correlation between a high rate response and the increasing trend of the one-way delay within a probe-packet train. Thus, Pathload is tied to the rate response curve. A change in the curve changes the estimates produced by Pathload.

In [11] a description on how the rate response curve affect estimates produced by Pathload as well as other available bandwidth methods (e.g. Spruce and PTR) is described in more detail. BART directly use the rate response curve in order to estimate the available bandwidth. Thus, the study on how the rate response curve change due to properties in wireless 802.11 networks is very important.

B. The testbed

The testbed used in this work consists of 9 computers running Linux, shown in Figure 2. The link speed for each link is shown in the figure. The links between $Xw1$, $Xw2$ and $R1$ are 802.11b wireless links (sharing the same channel) while the link between S and $R1$ either can be a 802.11b wireless link or a 100 Mbps wired link.

The bottleneck in the experiments is either the wireless 802.11b link or the link between $R1$ and $R2$, depending on whether the link between S and $R1$ is wired or wireless.

The cross traffic in the testbed is generated at the nodes $Xw1$, $Xw2$ and $X1$. The cross traffic consists of UDP packets and is generated by a modified version of *tg* [14]. The cross traffic is either constant bit rate (CBR), exponential or Pareto distributed (shape = 1.5). Further, the cross traffic consists of 60 (46% of the packets), 148 (11%), 500 (11%) and 1500 (32%) byte packets. This distribution of packet sizes originates from findings in [15].

C. Research question and experiments

In this paper we want to identify and explain properties associated with bandwidth measurements in wireless 802.11 networks.

The measurements have been performed mainly using DietTopp. We elaborate on the impact of probe packet size, the cross-traffic distribution and on the number of cross-traffic generators in the wireless network. We compare our results regarding available bandwidth to results obtained from Pathload.

The goal of the measurements performed is to understand the impact of the variables described above on the rate response curve. That is, the impact on the estimated available bandwidth (and link capacity) reported by many bandwidth measurement methods.

III. Experimental results

This section presents the results obtained from running DietTopp in different experiment scenarios. We have used Pathload to compare the obtained measurement results. In the diagrams all measurement results are shown with a 95% confidence interval.

A. Measurement results in wireless networks

This subsection presents the results from measurements using DietTopp and Pathload where the bottleneck is a wireless link (the link between S and R1 in the testbed as described in subsection II-B). Cross traffic is present on both of the wired links R1 - R2 and R2 - R3, but the rate is limited to approximately 9% of the corresponding link capacity (100 Mbps in this case). That is, the wireless link is the link that limits both the link capacity and the available bandwidth. The cross traffic at the 100 Mbps links between R1, R2 and R3 is Pareto distributed (with respect to cross-traffic packet arrival times) and consists of 4 different packet sizes. The cross-traffic configuration on the wired links is the same for each experiment presented in this section.

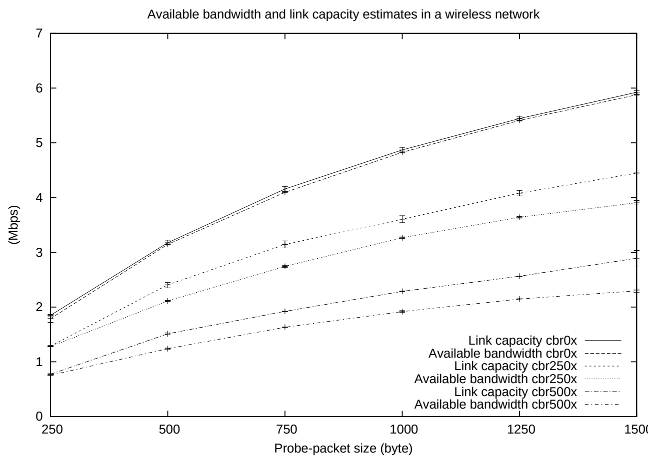

Fig. 3. Measured available bandwidth and link capacity using different probe-packet sizes under the impact of 0, 250 Kbps and 500 Kbps cross-traffic rates.

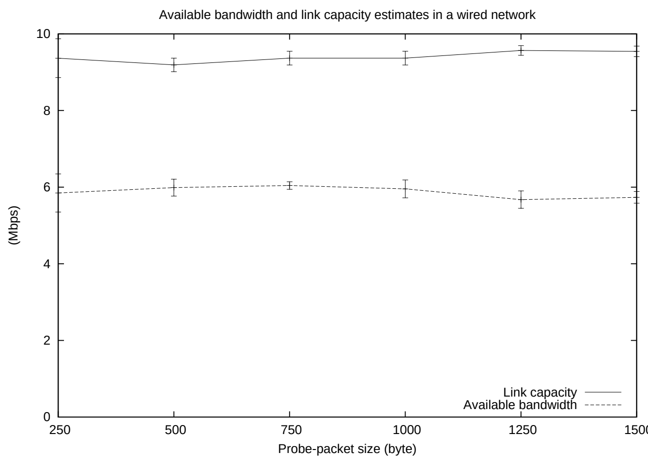

Fig. 4. Available bandwidth and link capacity measured by DietTopp in a wired network using different probe packet sizes. The cross traffic is a 3.26 Mbps Pareto distributed stream on a 10 Mbps link.

The probe-packet size affects both the measured link capacity and the available bandwidth estimate obtained by DietTopp when the bottleneck on an end-to-end path is a wireless 802.11 link. We illustrate and describe this phenomenon in a set of diagrams below.

The two upper curves in Figure 3 show the measured link capacity and the measured available bandwidth when no cross traffic is present on the wireless link. Varying the probe packet size from 250 bytes up to 1500 bytes gives increasing values of both the measured link capacity and the measured available bandwidth. It should be observed that the total number of bits remains constant independent of the probe packet size. The total amount of probe data sent by DietTopp in these measurements is 1.2 Mbit. Each probe train consists of 16 probe packets and we send 5 probe trains on each probe rate level. The number of probe rate levels depends on the probe packet size; increasing the probe packet size decreases the number of probe rate levels.

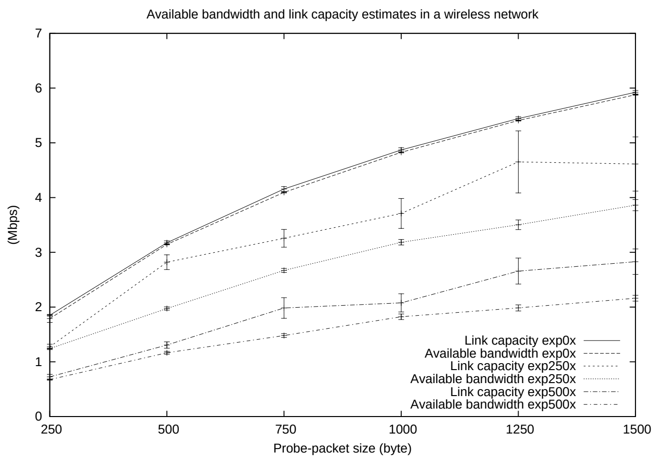

Fig. 5. Available bandwidth and measured link capacity measured under the impact of 0, 250 Kbps and 500 Kbps exponentially distributed cross-traffic.

Cross traffic	Measurement (Mbps)
0	2.32 - 2.39
250k cbr	1.67 - 1.67
250k exp	1.73 - 1.73
250k par	1.40 - 1.63
500k cbr	0.96 - 0.99
500k exp	0.87 - 0.95
500k par	1.27 - 1.29

TABLE I

MEASUREMENT RESULTS OBTAINED FROM RUNNING PATHLOAD IN DIFFERENT SCENARIOS WITH VARYING CROSS-TRAFFIC INTENSITY AND DISTRIBUTION.

The two middle curves show measurement estimates when there is a 250 Kbps constant-bit-rate (CBR) cross-traffic stream competing with the probe packets on the wireless link. The two bottom curves correspond to the case when a 500 Kbps CBR stream is present. *Both the measured link capacity and the measured available bandwidth increase with increasing probe-packet size*. Another equally important observation to be made is that *the measured link capacity decreases when increasing the cross-traffic rate*. Yet another interesting phenomenon is that the *difference between the measured link capacity and the measured available bandwidth tends to be smaller for small probe packet sizes*.

For comparison we have varied the probe packet size in an all wired network. That network is essentially the network shown in Figure 2 but the link between S and R1 is now a 100 Mbps wired link. The bottleneck is a 10 Mbps link between $R1$ and $R2$. The measurement results can be seen in Figure 4. Both the measured link capacity and the available bandwidth reported by DietTopp are relatively accurate and stable, that is independent of the probe packet size.

We have also performed measurements using Pathload, a tool that estimates the available bandwidth using 300 byte packets. The results obtained from executing Pathload in our testbed with different cross-traffic distributions and intensities

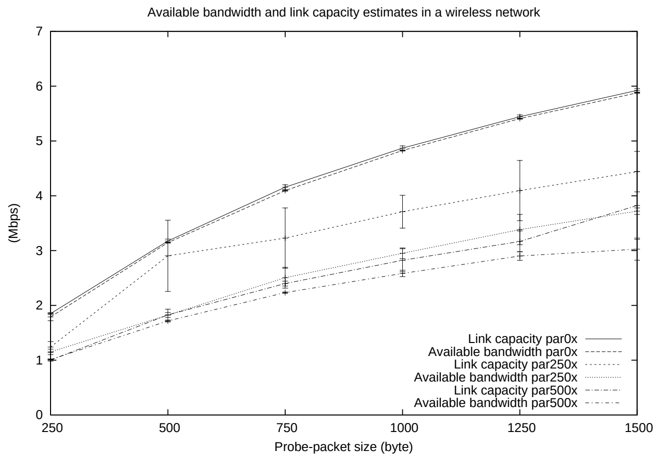

Fig. 6. Available bandwidth and measured link capacity measured under the impact of 0, 250 Kbps and 500 Kbps Pareto distributed cross-traffic.

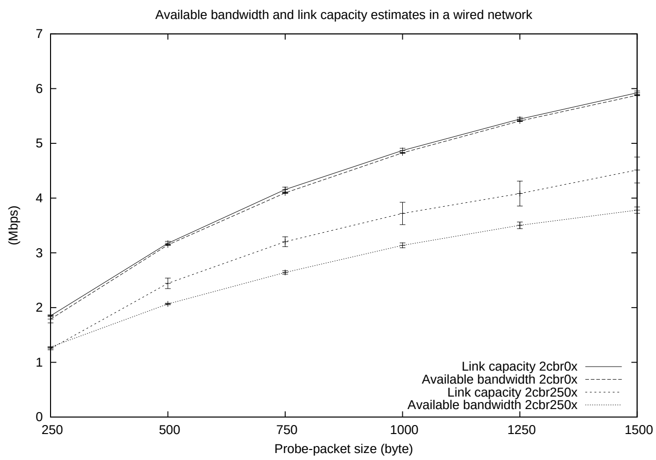

Fig. 7. Available bandwidth and measured link capacity measured under the impact of 0 and 500 Kbps CBR cross-traffic. The cross traffic is generated by two different sources that is injecting 250 Kbps each.

can be seen in Table I. When comparing results obtained by Pathload (in Figure 3) to those of DietTopp we can see that Pathload reports available bandwidth measurement estimations that are in line with estimations made by DietTopp (using interpolation between packet sizes 250 and 500 bytes).

Figures 5 and 6 report results from the same type of measurements that the ones Figure 3 is based on. The available bandwidth and the measured link capacity increases with increasing probe-packet size. Further, both the available bandwidth and the measured link capacity decrease with increasing cross-traffic rate. However, in these two scenarios more complex cross-traffic distributions are used. In Figure 5, exponentially distributed arrival times for the cross-traffic packets is used while in Figure 6 Pareto distributed arrival times is used. As can be seen in both figures, the confidence intervals are larger when the cross traffic is burstier. It is also obvious that the curves are less smooth compared to the CBR case in Figure 3. In the Pareto case (Figure 6) it is hard to distinguish between estimates of the link capacity in the two different cross-traffic scenarios (250 Kbps and 500 Kbps cross traffic). However, we can still see that the measured link capacity and available bandwidth are dependent on both the probe packet size and the cross-traffic rate. Again, comparing the measurement results (at the 300 byte probe packet size level) with results obtained by Pathload (in Table I) we can conclude that the available bandwidth estimate characteristics are compatible.

The somewhat blurred measurement results obtained when the cross traffic is more bursty (i.e. exponentially or Pareto distributed) can possibly originate from the fluid-like TOPP model. As mentioned above, a more complete, non-fluid, model has been developed in [11]. However, in the wired case which results are shown in Figure 4 no large fluctuations of the estimated bandwidth are visible, even though the cross traffic is Pareto distributed. Also, the estimates are rather accurate. The problem of the blurred measurement results is part of future research.

Continuing to Figure 7. In this scenario, two cross-traffic generators are generating 125 Kbps of CBR cross traffic each (that is summed up to 250 Kbps). As before, the probe-packet size affects both the measured link capacity and available bandwidth. Comparing Figure 7 to the measurement results shown in Figure 3 we see that the confidence intervals are larger when having multiple cross-traffic generators. Otherwise the curves in Figure 3 and in Figure 7 are similar. Hence, several cross-traffic generators on different nodes in a wireless network seem to increase the uncertainty, but not the overall estimates produced by DietTopp.

In the next section the results presented above will be discussed and analyzed in detail. This analysis result in an extended TOPP model for describing bandwidth measurements results obtained from networks with a wired or a wireless bottleneck.

IV. Analysis of experimental results

The reason for the varying measurement estimates of the link capacity and the available bandwidth in the experiments (as seen in the previous section) can be derived from the link-level acknowledgements and the contention phase used in 802.11 networks [16]. That is, if a probe packet is small, the relative overhead induced by the link-level acknowledgement and the contention phase is larger than if the probe packet were large. This will affect the probe-packet separation, which is the basis of the the rate response curve. Thus, the available bandwidth estimates produced by DietTopp, Pathload and other methods will vary with varying probe-packet size.

A. The extended TOPP model

The basic assumption in the TOPP model is, as described above, that the injected probe packets gain a proportional share of the link capacity when traversing a bottleneck link (if first-come-first-served is used as the queue policy). That is, the relationship during congestion (or overload) is described as

$$m = \frac{o}{(x+o)} * L \quad (3)$$

where m is the measured probe rate (at the receiver side of the network), o is the offered probe rate, x is the cross-traffic rate on the link and L is the link capacity on the specific link. Converting this equation into its linear form gives

$$\begin{aligned} \frac{o}{m} &= (1 - \frac{(L-x)}{L}) + \frac{o}{L} \quad (4) \\ &= \alpha * o + \beta. \end{aligned}$$

Now, assume that the link capacity can be rewritten as $L = s/T$ where s is the size in bits and T is the time it takes for the link to transmit s bits (assuming the queueing delay to be zero). Further, the offered rate o can be written as s/t_o where s again is the bit size and t_o the separation required in order to obtain a specific offered rate o. Then Equation 4 can be rewritten as

$$\frac{o}{m} = (1 - \frac{(s/T - x)}{s/T}) + \frac{s/t_o}{s/T} \quad (5)$$

Now, in a 802.11 wireless networks the time it takes to transmit one packet of size s over the link can be expressed as [16]

$$\begin{aligned} T &= T_{sifs} + T_{BO} + T_{difs} + T_{ack} + T_s \quad (6) \\ &= T_k + T_s \end{aligned}$$

where T_{sifs} and T_{difs} correspond to the time it takes to access the 802.11 wireless link, T_{BO} is the backoff time and T_{ack} is the time it takes for a link-level acknowledgment to return after sending a packet over the wireless link. T_{BO} is uniformly distributed in the interval $[0, CW-1]$ where CW grows with each unsuccessful transmission (due to competing cross traffic). Thus, T_k in the above equation, is an increasing function of the cross-traffic intensity x (i.e. $T_k = T_k(x)$).

Further, T_s is the time in seconds to transmit a packet of size s over the wireless channel. This time is determined by how the wireless signal is modulated, settings in the wireless network card and the size of the packet to be sent.

Substituting T in Equation 5 with the expression for T in Equation 6

$$\frac{o}{m} = \frac{T_k(x) + T_s}{s} * x + (T_k(x) + T_s)\frac{1}{t_o}. \quad (7)$$

Assume in the above equation that the probe-packet sender is injecting probe packets with a size of s bits. If changing the probe-packet size to s/n bits, where n is an arbitrary number ($n \geq 1$), the separation between the probe packets must also change in order to obtain the same rate o (i.e. t_o/n). Further, T_s is decreased to T_s/n. Thus

$$\begin{aligned} \frac{o}{m} &= \frac{T_k(x) + T_s/n}{s/n} * x + (T_k(x) + T_s/n)\frac{1}{t_o/n} \Longrightarrow \\ \frac{o}{m} &= \frac{n * T_k(x) + T_s}{s} * x + (n * T_k(x) + T_s)\frac{1}{t_o} \quad (8) \\ &= \beta + \alpha * \frac{1}{t_o}. \quad (9) \end{aligned}$$

From Equation 8, which describes the rate response curve when the bottleneck is a wireless 802.11 link, it can be seen that a decreasing probe-packet size (i.e. $n > 1$) increases both α and β in Equation 9. This, in turn decreases both the measured link capacity and the measured available bandwidth according to the TOPP model described in Section II-A. This behavior is visible in the diagrams in Section III that shows measurements in wireless 802.11 networks. Further, when increasing the cross-traffic intensity, $T_k(x)$ will increase due to increased contention and backoff time. This is also visible in the diagrams shown in the previous section where both the measured link capacity and the available bandwidth decrease with increasing cross-traffic intensity.

When the bottleneck is not a multi-access wireless link (i.e. when $T_k(x) = 0$) but rather a single access wired link the probe-packet size does not have impact on the estimates, which is seen in Equation 8. (T_s/s is constant.) The measurement results in Figure 4 support the above statement.

The important conclusion to be drawn from this section is that wireless 802.11 networks changes the parameters describing the rate response curve in Figure 1 when the probe-packet size is allowed to vary. A small probe-packet size increases the slope of b (i.e. decreases the estimate of the link capacity) while at the same time the intersection of b and $y = 1$ moves towards zero (i.e. decreases the estimate of the available bandwidth). This fact affects the estimates produced by DietTopp, Pathload and other methods that directly on indirectly rely on the same basic bandwidth estimation model.

The original TOPP model also describes how the relation between o and m changes when several congestable links (a link that will be congested due to the probe packets) between the sender and the receiver are present. The above supplement to the TOPP model can of course be incorporated into the case with several congestable links.

B. Estimated vs. physical link transmission capacity of a 802.11 wireless link

There is a need to differentiate between the estimates of the link capacity produced by DietTopp, that is not constant with respect to the cross-traffic rate and the probe-packet size, and the fixed physical link transmission capacity of the 802.11 wireless link. Using the equations in the previous section the differentiation is discussed in this section.

The TOPP model suggests that the bottleneck link capacity is computed as $1/\alpha$, where α is the slope of the straight line b shown in Figure 1. As derived in the previous section $\alpha = n * T_k(x) + T_s$ where T_s is the actual transmission time (which is constant if no change in the topology is made). $T_k(x)$

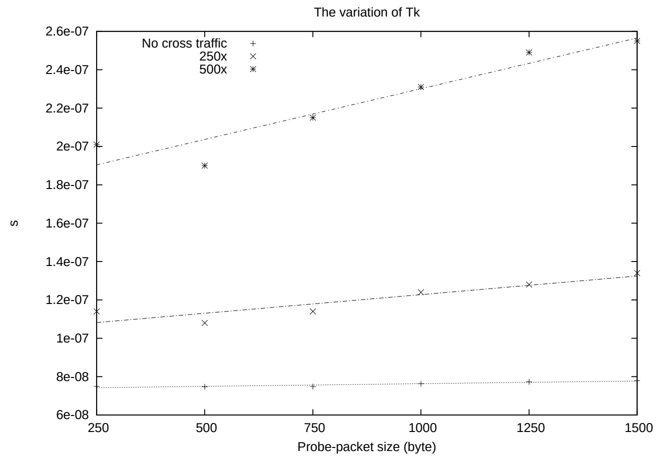

Fig. 8. The variation of T_k with increasing probe-packet size.

corresponds to contention and backoff time in multiple access networks, such as wireless 802.11. Thus, calculating $1/T_s$ gives the physical link transmission capacity of the wireless bottleneck link. First, assume that $T_k(x)$ is constant, with respect to the probe-packet size, then T_s is easily calculated by solving the equation system below

$$\begin{cases} \frac{1}{\alpha_1} &= n_1 * T_k(x) + T_s \\ \frac{1}{\alpha_2} &= n_2 * T_k(x) + T_s \end{cases} \tag{10}$$

where α_1 and α_2 correspond to link capacity estimates while n_1 and n_2 are fixed during the measurement.

In Figure 3 (CBR case) the estimated link capacity is 5.92 Mbps ($= \alpha_1$) when no cross traffic is present using 1500 byte probe packets (fix $n_1 = 1$). Further, the estimated link capacity is approximately 4.16 Mbps ($= \alpha_2$) in the same scenario but when using 750 byte probe packets (then $n_2 = 2$). That is, $n_1 = 1$ and $n_2 = 2$ since we use half the probe-packet size is the second scenario. Solving the above equation system with these values as input $T_s = 8.34^{-8}$ which in turn gives a physical link transmission capacity of 10.5 Mbps. This estimate is fairly close to 11 Mbps which is the link transmission capacity in the experiments.

However, solving similar equation systems using input from scenarios with competing cross traffic, T_s will not correspond to a link transmission capacity of 11 Mbps at all. [1] This is impossible, since T_s is a physical, constant property of the wireless link (in our experiments). Thus, T_k must vary with the probe-packet size for some reason. In Figure 8 the variation of T_k is shown, assuming T_s to be constant ($T_s = 9.09^{-8}$ which corresponds to a link transmission capacity of exactly 11 Mbps). As can be seen, the variation of T_k can be described by a straight line.

Further, as seen in the figure, T_k increases with increasing cross-traffic intensity independent of the probe-packet size. This is obvious since a high cross-traffic intensity increases the backoff time T_{BO} in Equation 6. It should also be noted that the slope of T_k increases with increasing cross-traffic intensity.

We believe that the observed dependency of T_k on the probe-packet size can be traced to the link-layer retransmission mechanism in 802.11 wireless networks [16]. A 802.11 wireless node cannot tell whether a collision on the link has occurred or not since the radio used cannot send and listen at the same time. That is, the only way to detect a collision (or other link-layer errors that requires a retransmission) is the lack of a link-layer acknowledgment that should be transmitted back by the receiver node. When no link-layer acknowledgment is received, the link-layer on the sending side retransmits the packet. The question is; why is the packet size crucial to T_k when a link-layer retransmission is triggered? There are probably several reasons for this and we elaborate on two below.

Collisions: In the case corresponding to the measurement results presented in this work, the number of transmitted bits (1.2 Mbit) is constant independently of the probe-packet size, as described in Section III. Hence, sending one 1500 byte probe packet corresponds to sending six 250 byte probe packets. If one collision occur, in either case, one probe packet has to be retransmitted by the link layer. The procentual increase in T_k will be much higher using 1500 byte probe packets compared to using 250 byte probe packets. That is why T_k increases with increasing packet size.

Noisy wireless link: If the wireless link is noisy the probability for bit errors in the transmitted packet increases. Assuming that a constant number of probe packets are sent, the probability for encountering bit errors will increase with the probe-packet size. That is, the probe packet must be retransmitted and T_k increases.

It seems that T_k is actually a function of both the cross traffic and the packet size. That is, $T_k = T_k(x, s)$. Assuming the cross-traffic intensity is constant then $T_k(s) = \phi + \varphi * s$ where s is the probe-packet size while ϕ and φ are constants describing the suggested straight lines in Figure 8.

To be able to calculate T_s, that is the value corresponding to the physical link transmission capacity of the wireless link the function $T_k(x, s)$ must be known. In the case where no cross traffic flows on the wireless link, T_k is rather constant (see Figure 8) and thus, T_s could be calculated, as shown above. However, whether it is possible to calculate T_s in the general cross-traffic intensity case is left for future research.

V. Other observations

Due to the fact that the probe-packet size affects the rate response curve shown in Figure 1 (and thus affects both the measured link capacity and the measured available bandwidth when using DietTopp), a possible method to identify a wireless bottleneck in a network path could be: if the available bandwidth (and the measured link capacity) changes when probing the path with different probe-packet sizes, this can be taken as an indication that the path includes a wireless bottleneck. This is important since, as we have discussed, wireless bottlenecks have different characteristics than wired bottlenecks.

[1] The exact numbers for the estimated link capacity is left out, but a rough approximation can be obtained from the figure.

An important consequence of the measurements we have presented in this paper is that the available bandwidth will be application dependent if the bottleneck is a wireless link. For example, a voice over IP application or a distributed game probably use small packets while a file transfer may use large packets. The available bandwidth for the applications will not be the same due to their packet-size distribution. This means that when probing a path containing a wireless bottleneck link, the estimation tool must use a probe-packet size distribution that corresponds to the specific application that is to use the estimates.

VI. Conclusion

In this paper we have shown measurement results that illustrate the difference between available bandwidth measurement results obtained in wired and wireless networks. The impact of the probe-packet size, the cross-traffic rate and the number of cross-traffic generators has been analyzed. In the paper, we have discussed in detail the underlying reasons for these differences by extending the TOPP model. The extended model shows that the rate response curve has different properties when performing active end-to-end bandwidth measurements in 802.11 networks compared to the one obtained in ordinary wired networks.

Using the extended model it has been shown why the probe-packet size, for example, is crucial to the available bandwidth estimate. Further, we have shown why the estimated link capacity reported by TOPP is erroneous when performing measurements where the bottleneck is a wireless 802.11 link. A very important aspect of the extended TOPP model is that it also have implications on the estimates reported by other end-to-end available bandwidth measurement methods (e.g. BART, DietTopp, Pathload and PTR).

The estimates of the available bandwidth and link capacity have been produced using our own tool, DietTopp. For comparison and validity we have used Pathload. The measurements have been performed in a testbed where we have used different types of cross traffic, from simple CBR to bursty Pareto distributed cross traffic.

Future research is to study whether the physical link capacity of 802.11 wireless bottleneck links can be estimated in general scenarios using DietTopp. Another open issue is to study how the 802.11 properties can fit into a non-fluid model for describing available bandwidth measurement methods. It is also important to investigate why the use of small probe packets results in a lower variance when used for active probing in wireless 802.11 networks.

References

[1] S. Ekelin, M. Nilsson, E. Hartikainen, A. Johnsson, J. Mångs, B. Melander, and M. Björkman, "Real-time measurement of end-to-end available bandwidth using kalman filtering," in *upcoming Proceedings to the Network Operations and Management Symposium*, Vanconver, Canada, 2006.

[2] Ribeiro, Riedi, Baraniuk, Navratil, and Cottrell, "pathchirp: Efficient available bandwidth estimation for network paths," in *Passive and Active Measurement Workshop*, 2003.

[3] Manish Jain and Constantinos Dovrolis, "End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput," in *Proceedings of ACM SIGCOMM*, Pittsburg, PA, USA, Aug. 2002.

[4] Ningning Hu and Peter Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," *IEEE Journal on Selected Areas in Communication*, 2003.

[5] Strauss, Katabi, and Kaashoek, "A measurement study of available bandwidth estimation tools," in *ACM SIGCOMM Internet Measurement Workshop*, 2003.

[6] Bob Melander, Mats Björkman, and Per Gunningberg, "Regression-based available bandwidth measurements," in *Proceedings of the 2002 International Symposium on Performance Evaluation of Computer and Telecommunications Systems*, San Diego, CA, USA, July 2002.

[7] Andreas Johnsson and Bob Melander Mats Björkman, "Bandwidth measurement in wireless networks," in *Proceedings to the Fourth Annual Mediterranean Ad Hoc Networking Workshop*, Porquerolles, France, 2005.

[8] R.S. Prasad, M. Murray, C. Dovrolis, and K. Claffy, "Bandwidth estimation: metrics, measurement techniques, and tools," IEEE Network Magazine, 2003.

[9] Andreas Johnsson, Mats Björkman, and Bob Melander, "A study of dispersion-based measurement methods in ieee 802.11 ad-hoc networks," in *Proceedings of the International Conference on Communication in Computing*, Las Vegas, 2004.

[10] Karthik Lakshminarayanan, Venkata N. Padmanabhan, and Jitendra Padhye, "Bandwidth estimation in broadband access networks," in *In Proceedings to the Internet Measurement Conference*, 2004.

[11] Xiliang Liu, Kaliappa Ravindran, and Dmitri Loguinov, "Multi-hop probing asymptotics in available bandwidth estimation: Stochastic analysis," in *Proceedings of ACM SIGCOMM Internet Measurement Conference*, 2005.

[12] Andreas Johnsson, "Diettopp implementation, beta version," http://www.idt.mdh.se/ãjn12/, 2005.

[13] Constantinos Dovrolis, Parameswaran Ramanathan, and David Moore, "What do packet dispersion techniques measure?," in *Proceedings of IEEE INFOCOM*, Anchorage, AK, USA, Apr. 2001, pp. 905–914.

[14] Paul E. McKenney, Dan Y. Lee, and Barbara A. Denny, "Traffic generator software release notes," SRI International and USC/ISI Postel Center for Experimental Networking, Jan. 2002.

[15] Yongmin Choi, Heung-No Lee, and Anurag Garg, "Measurement and analysis of wide area network (wan) traffic," in *SCS Symposium on Performance Evaluation of Computer and Telecommunication Systems*, July 2000.

[16] "Ieee 802.11," http://grouper.ieee.org/groups/802/11/.