Bandwidth Measurement in Wireless Networks

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

The Department of Computer Science and Engineering
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, such as ad-hoc networks, to be no different than the one in wired networks. However, in networks with 802.11 wireless links we show that this is not the case. We also discuss the underlying reasons for the observed differences.**

Our experiments 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 is dependent on the probe packet size *and* on the cross-traffic intensity.

The study we present has been performed using a bandwidth measurement tool, DietTopp, that is based on the previously not implemented TOPP method. DietTopp measures the end-to-end available bandwidth of a network path along with the capacity of the congested link.

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 performance measurements in terms of network bandwidth, both link bandwidth and the unused portion thereof; the available bandwidth.

Measurement of network properties such as available bandwidth in for example ad-hoc 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. Since the exact route between two nodes in an ad-hoc network usually is unknown and may change without notification to the application layer the end-to-end measurement of the available bandwidth should not require any infrastructure or pre-installed components at each node. To achieve that, common end-to-end bandwidth measurement methods can be applied.

State-of-the-art bandwidth measurement methods are for example Pathchirp [1], Pathload [2], Spruce [3] and TOPP [4]. 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 packets. By analyzing these time stamps estimates of the link capacity and/or the available bandwidth can be made. For many end-to-end available bandwidth measurement methods no previous knowledge of the underlying network topology is needed. That is, bandwidth estimation methods are well suited for end-to-end performance measurements in ad-hoc networks.

The existing methods differ in how probe packet are sent (the flight patterns) and in the estimation algorithms used. An overview of methods and tools in this area can be found in [5].

In the following sections, we describe and measure bandwidth estimation characteristics when probing in 802.11 wireless networks. We show that both the measured available bandwidth and the measured link capacity are dependent on the probe packet size. Furthermore, our measurements indicate that the measured link capacity is also dependent on the cross-traffic rate. We discuss the origins of some of the observed behavior.

The measurements have been performed in a testbed containing both wireless and wired hops. Our testbed topology only consist of one wireless hop, but we believe that our results illustrate the measurement problem for larger ad-hoc networks, consisting of several wireless hops, as well. To produce measurement results we have used DietTopp, a tool that measures the available bandwidth and link capacity of an end-to-end path. For comparisons and to illustrate that our observations are not tied to a certain measurement tool, we have also used the tool Pathload, that measure the available bandwidth of an end-to-end path, in our experiments.

Earlier work has touched upon the problem of active measurements of bandwidth in wireless networks. In [6] we discuss the main problem areas when deploying existing bandwidth measurement methods in ad-hoc networks. For example, we observed using ns-2 simulations, that the measured link capacity show dependence on the cross-traffic rate.

Measurement results presented in [7] indicate that the available bandwidth is dependent of the probe packet size. Our study extends that study by showing that *both* the available bandwidth and the measured link capacity depends on both the probe packet size and the cross-traffic rate. Further, we use a more complex measurement topology to verify their findings.

The rest of this paper is organized as follows. Section II-A describes the original TOPP measurement method. DietTopp, which is our implementation of a modified TOPP method, 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. Section III shows 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 some important observations are made. The paper ends with conclusions in Section V.

II. EXPERIMENTAL SETUP

This section describes our experimental setup. That is, the measurement tool (DietTopp), our testbed and what kind of measurements we have performed and their relevance to ad-hoc networks.

A. DietTopp

DietTopp has its origins in the previously not implemented TOPP [4] method and uses the measured dispersion of probe packet trains to calculate bandwidth estimates.

In short summary DietTopp works as follows. Starting at some offered probe rate o_{min}, DietTopp injects m probe packet trains, where each train contains k equally sized probe packets, into the network path. When all probe trains corresponding to a probe rate o_{min} have been transmitted, DietTopp increases the offered rate o by Δo. Another set of probe packet trains are sent into the network with the new probe rate. This is repeated i times until the offered probe rate reaches some specified probe rate o_{max}.

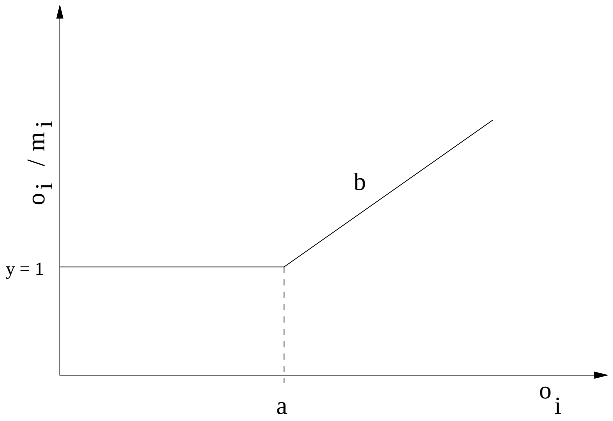

Fig. 1. Plot of the ratio o_i/m_i as a function of o_i.

The probe packet dispersion may change as the probe packets traverse the network path between the probe sender and the probe receiver. This is due to the *bottleneck spacing effect* [9] and/or interactions with competing traffic.

The receiver time stamps each probe packet arrival. Hence, any change in probe packet separation can be measured. The time stamps are used to calculate the measured probe rate m_i.

When all measurements are collected, DietTopp computes the ratio o_i/m_i for all i. If plotting the ratio o_i/m_i on the y-axis and o_i on the x-axis for all i, we get a plot like the theoretical one in Figure 1. If the dispersion of the probe packets would remain unchanged after traversal of the network path, the measured rates, m_i, on the receiver side would be the same as the offered rates o_i. Expressed differently, the ratio o_i/m_i would equal 1. The link that limits the available bandwidth of the path will eventually get congested when increasing the offered probe rate. This causes the curve to rise since the rate m does not increase as much as the rate o. If the link capacity is l and the available bandwidth is a the relation between o_i and m_i is given by $o/m = (1 - a/l) + o/l$ (when one link is congested) [4].

Segment b in the figure is linear and the slope corresponds to the link capacity of the congested link. The available bandwidth of the end-to-end path is defined as the intersection of $y = 1$ and b (i.e. a in the figure) [4].

To speed up the probing phase of DietTopp we want to avoid measurements below a. That is, we want 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} and then measure their separation at the receiver. According to [4] m_{max} is greater than the available bandwidth (m_{max} is referred to as the asymptotic dispersion rate in [10]).

Having a value of $o_{min} > a$ the procedure described above is executed to find the link capacity and available bandwidth.

DietTopp is implemented in C++ on Unix platforms and can be downloaded from [11].

B. The testbed

The testbed used 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 while the link between S and $R1$ either can be a 802.11b wireless link or a 100 Mbps wired link.

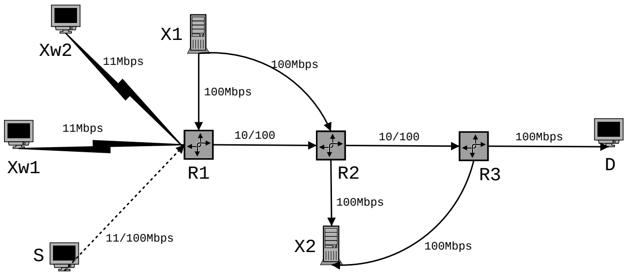

Fig. 2. The testbed is constructed by one wireless link, three routers and several cross-traffic generators (on both the wireless and the wired links)

The cross traffic, generated by a modified version of *tg* [12], can either take the route $X1 \rightarrow R1 \rightarrow R2 \rightarrow X2$ or the route $X1 \rightarrow R2 \rightarrow R3 \rightarrow X2$. Cross traffic can also be generated by *Xw1* and *Xw2* on the wireless hop. The cross traffic is either constant bit rate (CBR), exponential or pareto distributed (shape = 1.5). Further, the cross traffic consists of 60 (46%), 148 (11%), 500 (11%) and 1500 (32%) byte packets. This distribution of packet sizes originates from findings in [13].

C. Experiments

In this paper we want to identify possible problems associated with bandwidth measurements in wireless networks, such as ad-hoc networks. First we show two measurements using DietTopp in a wired scenario. This is to validate that our tool is sound in the simple wired case before turning attention to the more complex case of estimating end-to-end bandwidth in wireless networks. We compare DietTopp results to theoretical values as well as to values obtained from Pathload.

The measurements in the wireless scenario is done using DietTopp. We elaborate on the impact of probe packet size, the cross-traffic distribution, the number of probe packets sent and on the number of cross-traffic generators in the wireless network. We compare our results to results obtained from Pathload.

This work is related to the work presented in [7]. We extend and complement that work in the following way: We use our newly developed tool DietTopp, that measured both the link capacity and the available bandwidth of the bottleneck link. Previous work has only focused on the available bandwidth on wireless links. Further, we use a more complex testbed topology.

III. EXPERIMENTAL RESULTS

This section presents the results obtained using DietTopp in wired and wireless scenarios. We have used Pathload [2] to compare and discuss the obtained measurement results. In the diagrams all measurement results are shown with a 95% confidence interval.

A. Measurement results in wired networks

This section presents measurements done with both DietTopp and Pathload in an all wired scenario. This section is to show by example that our tool, DietTopp, measures both the link capacity and the available bandwidth in a sound way.

The diagram in Figure 3 illustrates results from DietTopp measurements using four different cross traffic intensities on link R1 - R2 (10 Mbps link capacity in this case), shown on the x-axis. The cross traffic at link R2 - R3 (100 Mbps link capacity) is a 8.76 Mbps stream. Both cross-traffic streams are exponentially distributed. The y-axis shows the measured link capacity (thin solid line), the measured available bandwidth (thin dashed line), the theoretical link capacity (thick solid line) and the theoretical available bandwidth (thick dashed line). As can be seen the correlation between measurement results and the theoretical values is good.

The diagram in Figure 4 is a comparison of the measured available bandwidth using DietTopp (dashed line) and Pathload (solid line). The same testbed and cross traffic setup is used as in Figure 3. We see that both tools report similar estimates of the available bandwidth.

We have now given an indication that DietTopp estimates both the link capacity as well as the available bandwidth in wired network with good accuracy, both compared to theoretical values and compared to one state-of-the-art bandwidth measurement tool, Pathload. In the next subsection we investigate the impact of wireless bottlenecks on the measurement results.

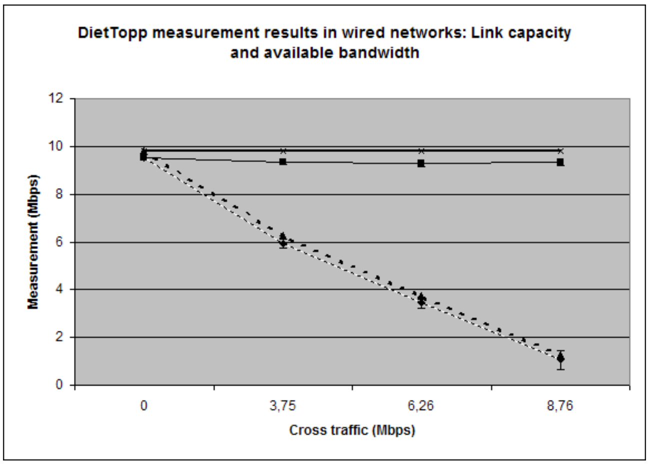

Fig. 3. Link capacity (solid lines) and available bandwidth (dashed lines). Thick lines corresponds to theoretical values while thin lines are values obtained from DietTopp.

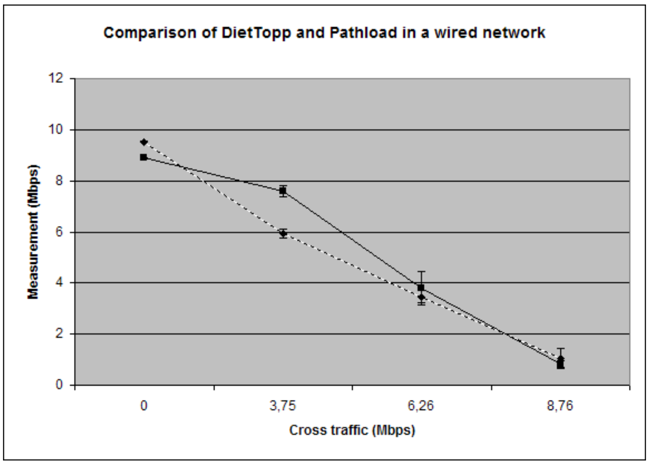

Fig. 4. Available bandwidth measured by DietTopp (dashed line) and Pathload (solid line).

B. *Measurement results in wireless networks*

This subsection presents our results from measurements using DietTopp where the bottleneck is a wireless link (the link between S and R1 in the testbed as described in subsection II-B) which is the case in ad-hoc wireless networks. 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 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.

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

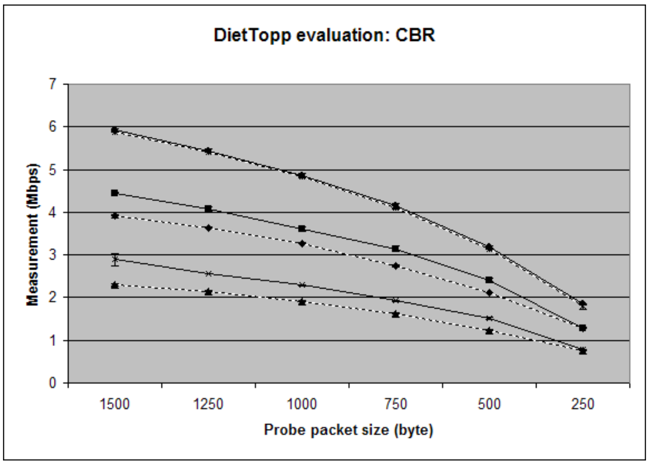

Fig. 5. Available bandwidth (dashed lines) and measured link capacity (solid lines) measured under 0, 250 Kbps and 500 Kbps cross-traffic rates).

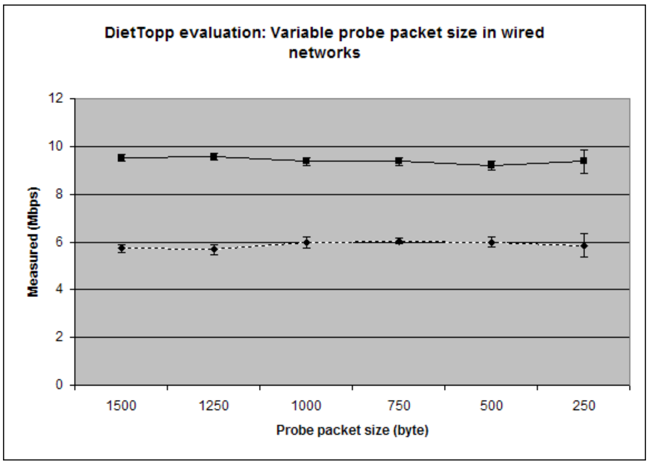

Fig. 6. Available bandwidth (dashed line) and link capacity (solid line) measured by DietTopp in a wired network using different probe packet sizes. The cross traffic is a 3.26 Mbps pareto distributed stream on a 10 Mbps link.

The two upper curves in Figure 5 show the measured link capacity (solid line) and the measured available bandwidth (dashed line) when no cross traffic is present on the wireless link. Varying the probe packet size from 1500 bytes down to 250 bytes gives decreasing 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

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 PATHLOAD UNDER THE INFLUENCE OF DIFFERENT CROSS-TRAFFIC DISTRIBUTIONS.

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; decreasing the probe packet size increases the number of probe rate levels.

The two middle curves show measurement values when there is a 250 Kbps CBR cross-traffic stream on the wireless link. The two bottom curves correspond to the case when a 500 Kbps CBR stream is present. The same decreasing trend for the measured link capacity and the measured available bandwidth is visible. An 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. Why this is the case is a subject of further research.

For comparison we have varied the probe packet size in an all wired network. The measurement results can be seen in Figure 6. Both the measured link capacity and the available bandwidth are quite stabile, that is independent of the probe packet size.

We have also done measurements using Pathload, a tool that estimates the available bandwidth using 300 byte packets. The results obtained from using Pathload in our testbed with different cross-traffic distributions and intensities can be seen in Table I. When comparing results obtained by Pathload (in Figure 5) 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 7 and 8 report results from the same type of measurements as in Figure 5. However, in these two scenarios we have used more complex cross-traffic distributions. In Figure 7 we have used exponentially distributed arrival times for the cross-traffic packets while in Figure 8 we have used pareto distributed arrival times. 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 5. In the pareto case (Figure 8) it is hard to distinguish between the 250 Kbps and 500 Kbps measurements of link capacity and available bandwidth. However, we can still see that the measured link capacity and available bandwidth is 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.

In Figure 9 we vary the probe packet size in the same

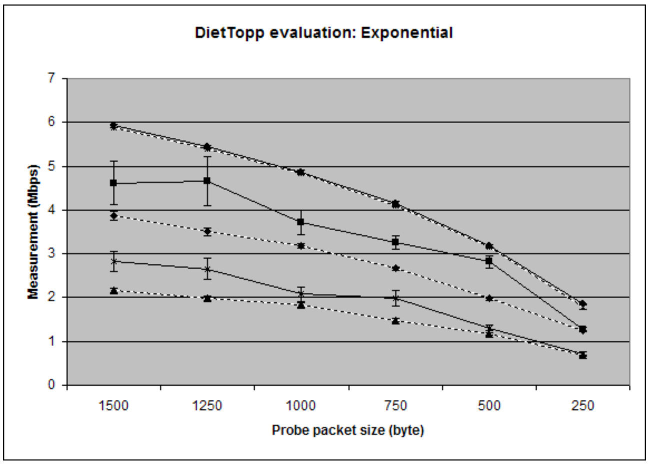

Fig. 7. Available bandwidth (dashed lines) and measured link capacity (solid lines) measured under 0, 250 Kbps and 500 Kbps exponentially distributed coss-traffic.

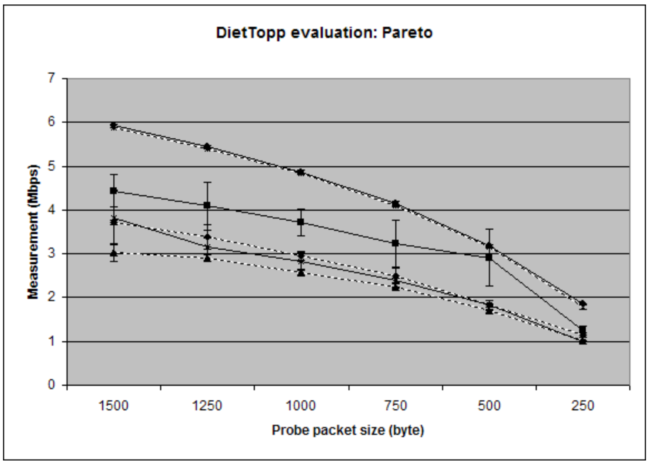

Fig. 8. Available bandwidth (dashed lines) and measured link capacity (solid lines) measured under 0, 250 Kbps and 500 Kbps pareto distributed cross-traffic.

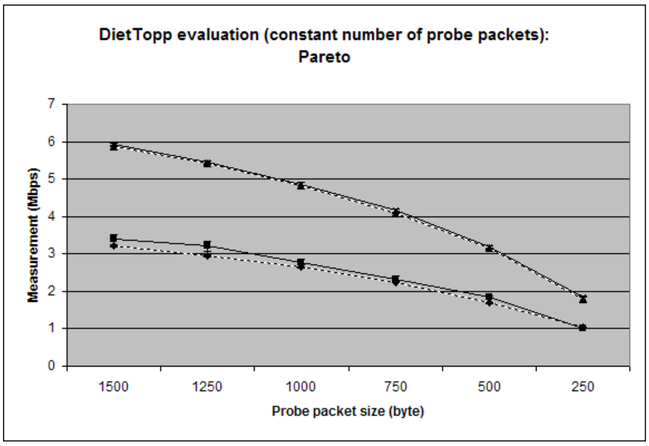

Fig. 9. Available bandwidth (dashed lines) and measured link capacity (solid lines) measured under 0 and 500 Kbps pareto distributed cross-traffic. The number of probe packets is constant.

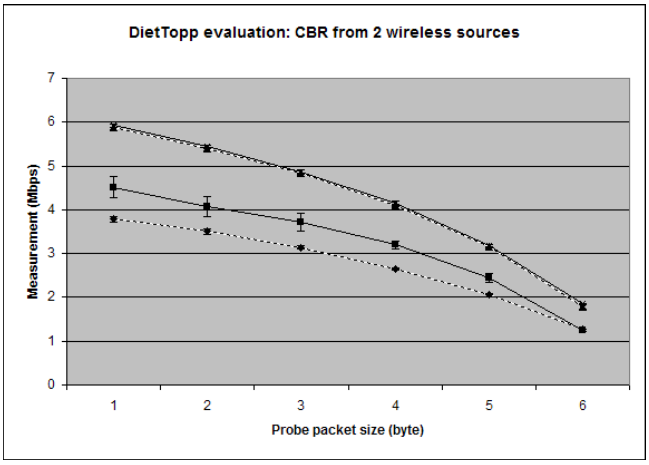

Fig. 10. Available bandwidth (dashed lines) and measured link capacity (solid lines) measured under 0 and 500 Kbps CBR cross-traffic. The cross traffic is generated by two different sources (250 Kbps each).

manner as above. However, instead of keeping the total number of bits transfered constant we keep the number of probe packets sent constant. The cross traffic is pareto distributed. We see that even though the total amount of probe data sent is less at each probe packet size level the confidence intervals remain low.

In Figure 10 two cross-traffic generators are generating 250 Kbps of CBR cross traffic each. Comparing Figure 10 to the measurement results in Figure 5 we see that the confidence intervals are larger when having multiple cross-traffic generators.

C. *Wireless measurement results discussed*

In this subsection we will discuss the results obtained in the previous subsection and the reasons for the difference between DietTopp measurements in wired and in wireless networks.

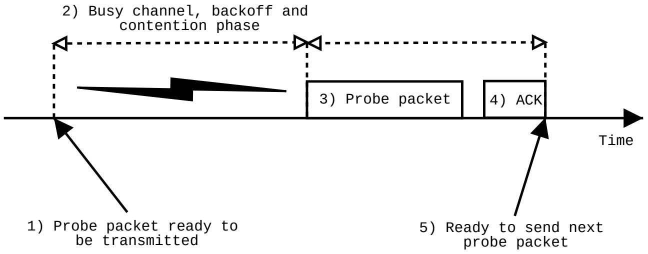

Fig. 11. A schematic picture of the procedure for sending a packet in a 802.11 wireless network.

We will derive the differences from Figure 11 which illustrates the procedure for sending a packet in a 802.11 wireless network. First, the radio transmitter at the wireless node needs a clear channel to send its packet on. This is illustrated by step 1 and 2 in the figure. If someone else is using the channel the sender does a back-off. It tries again after some time. Eventually the packet is sent, step 3 in the figure. When the receiving node gets the whole packet it responds with a link-layer acknowledgment to the sender (step 4). The sender can now transmit the next packet.

The reason for the decreasing measurement values of the available bandwidth can be derived from the link-level acknowledgments in step 3 and 4 in the figure. That is, if the probe packet is small, the overhead induced by the link-level acknowledgment is larger than if the probe packet were large. We come to the conclusion that large probe packets will measure a larger available bandwidth than small probe packets.

The contention phase (step 1 and 2 in the figure) is independent of the packet size. The contention phase is instead dependent on the number of sending nodes in the wireless networks. Increasing the number of stations that want to send traffic over the wireless network increases the waiting time for each node. It also increases the variance of the waiting time.

In Figure 10 two cross-traffic generators are generating 250 Kbps of CBR cross traffic each as described above. Since we have two wireless nodes sending traffic, this is likely to affect the contention phase in Figure 11 in such a way that we get larger confidence intervals in our measurement results. Comparing Figure 10 to the measurement results in Figure 5 we see that the confidence intervals are larger when having multiple cross-traffic generators.

The results concerning the available bandwidth are in line with results discussed in [7]. We validate and extend thouse findings by using more complex testbed scenarios and our own tool DietTopp.

A theoretical description of why the measured link capacity is dependent on both the probe packet size and the cross-traffic intensity is a subject of future research.

A final remark is that in most figures we can see that the confidence intervals decrease with the probe packet size. Hence, we can draw the conclusion that we get values with low standard deviation with small probe packets. However, why this is the case is also a subject of future research.

IV. OTHER OBSERVATIONS

Due to the fact that the probe packet size affects both the measured link capacity and the measured available bandwidth when using DietTopp, a possible method to identify a wireless bottleneck link in a network path could be: if the available bandwidth (and the measured link capacity) changes when probing the path with different 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. This is also interesting from an semi-ad-hoc perspective: when one node of an ad-hoc network is connected to an infrastructure, such as the Internet, it is important to determine whether the bottleneck is within the ad-hoc network or within the infrastructure. Is the bottleneck within the ad-hoc network there might be possibilities to route the data differently. Also, ad-hoc router protocols can perform better with an understanding of bottlenecks within the ad-hoc network. However, this subject is left to future research.

An important consequence of the measurements we have presented in this paper is that the available bandwidth will be application dependent in ad-hoc networks and when wireless links are a bottleneck in general. For example, a voice over IP application or a distributed game probably use small packets to send data while a file transfer application may use larger packets. The available bandwidth for the applications will not be the same due to their packet size distribution, as indicated by the figures above that show decreasing measurement values when decreasing the probe packet size. 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.

V. CONCLUSION

In this paper we have shown measurements that illustrate the difference between bandwidth measurements in wired and wireless networks, such as ad-hoc networks. We have discussed some of the underlying reasons for these differences while other reasons are left to further research. We have used our own tool, DietTopp, to produce measurement results throughout the paper. For comparison and validity we have used Pathload. The measurements have been performed in a testbed where we have used different kinds of cross traffic, from simple CBR to bursty pareto distributed cross traffic.

Our conclusions are that measurements in wireless networks are associated with difficulties that can result in misleading bandwidth estimations. We have shown that the packet size is critical to the bandwidth measurement value of both the link capacity and the available bandwidth. Further, we have shown that the measured link capacity on wireless links does not only depend on the packet size, but also on the cross traffic intensity. We have also addressed the problem of application dependent probing.

Future research is to investigate why small packets gives a lower variance when used for active probing in wireless networks. We will also investigate why the measured link capacity vary when the probe packet size vary. It is also important to study what the variable measured link capacity obtained means for wireless network applications.

REFERENCES

[1] Ribeiro, Riedi, Baraniuk, Navratil, and Cottrel, “pathchirp: Efficient available bandwidth estimation for network paths,” in *Passive and Active Measurement Workshop*, 2003.

[2] M. Jain and C. Dovrolis, “End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput,” in *Proceedings of ACM SIGCOMM*, Pittsburg, PA, USA, Aug. 2002.

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

[4] B. Melander, M. Björkman, and P. 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.

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

[6] A. Johnsson, M. Björkman, and B. 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.

[7] K. Lakshminarayanan, V. N. Padmanabhan, and J. padhye, “Bandwidth estimation in broadband access networks,” in *In Proceedings to the Internet Measurement Conference*, 2004.

[8] A. Johnsson, B. Melander, and M. Björkman, “Diettopp: A first implementation and evaluation of a new bandwidth measurement tool,” in *Swedish National Computer Networking Workshop*, 2004.

[9] V. Jacobson, “Congestion avoidance and control,” in *Proceedings of ACM SIGCOMM*, Stanford, CA, USA, Aug. 1988, pp. 314–329.

[10] C. Dovrolis, P. Ramanathan, and D. Moore, “What do packet dispersion techniques measure?” in *Proceedings of IEEE INFOCOM*, Anchorage, AK, USA, Apr. 2001, pp. 905–914.

[11] A. Johnsson, “Diettopp implementation, beta version,” http://www.idt.mdh.se/ãjn12/.

[12] P. E. McKenney, D. Y. Lee, and B. A. Denny, “Traffic generator software release notes,” SRI International and USC/ISI Postel Center for Experimental Networking, Jan. 2002.

[13] Y. Choi, H.-N. Lee, and A. Garg, “Measurement and analysis of wide area network (wan) traffic,” in *SCS Symposium on Performance Evaluation of Computer and Telecommunication Systems*, July 2000.