

Experimental Comparison of Bandwidth Estimation Tools for Wireless Mesh Networks

Dhruv Gupta Daniel Wu Prasant Mohapatra Chen-Nee Chuah
Graduate Group in Computer Science
University of California Davis
{dhgupta, danwu, pmohapatra, chuah}@ucdavis.edu

Abstract—Measurement of available bandwidth in a network has always been a topic of great interest. This knowledge can be applied to a wide variety of applications and can be instrumental in providing Quality of Service to end users. Several probe-based tools have been proposed to measure available bandwidth in wired networks. However, the performance of these tools in the realm of wireless networks has not been evaluated extensively. In recent years, there has also been some work on estimating bandwidth in wireless networks via passively monitoring the channel and determining the ‘busy’ and ‘idle’ periods. However, such techniques have primarily been evaluated via simulations only. In this work, we perform an extensive experimental comparison study of both passive and active bandwidth estimation tools for 802.11-based wireless mesh networks. We investigate the impact of interference, packet loss, and 802.11 rate-adaptation, on the performance of these tools. Our results indicate that for wireless networks, a passive technique provides much greater accuracy than the probe-based tools.

I. INTRODUCTION

At the physical layer, bandwidth refers to the width of the communication spectrum available for data transmission. However, at the network and higher layers, it refers to the data rate available on a network path or link, usually expressed in bits per second. The concept of bandwidth is very important for present day networks, and various applications can benefit from this information. The term bandwidth is often confused with different throughput related concepts. Authors in [1] provide an overview of how to differentiate between these terminologies. The *available bandwidth* of a link relates to the *unused or spare capacity* of the link during a given time period. It can be determined by finding the time period for which the link is not utilized for transmitting data. IEEE 802.11-based wireless mesh networks (WMNs) are being widely deployed, especially in the municipal and enterprise setting. In such scenarios, providing performance guarantees to end users may be an important goal for the network administrators or ISPs. Admission control and routing schemes based on available bandwidth (Av-Bw) can be very effective for provisioning Quality of Service (QoS) to end users. However, WMNs present several challenges to the bandwidth estimation process, such as interference, contention-based MAC, and rate-adaptation feature of 802.11 protocols.

While the performance of probe-based tools has been widely evaluated for estimating Av-Bw in wired networks, similar research is lacking for WMNs. Even though it has been widely cited that these tools do not work well for wireless networks [2] [3], their performance has not been quantified via detailed

experiments. Several recent works have also proposed using a passive approach for bandwidth estimation in wireless networks. However, these approaches have mostly been tested via simulations only and lack an extensive experimental evaluation. The objective of our study is to evaluate the existing tools and study their performance for 802.11-based WMNs. We also implement a passive tool that can estimate the channel utilization by continuously sniffing the wireless channel. We evaluate these tools by varying various parameters such as the physical data rate, number of hops, and amount of interference. Via extensive experiments, we show that the active probe-based tools do not perform well in wireless networks. The contention-based MAC, and the tendency of 802.11 protocol to provide fairness, cause these tools to provide inaccurate estimates. Our results indicate that the passive scheme provides accurate estimates even in the face of above mentioned challenges. Our main contributions are highlighted below:

- We select four probe-based tools, and one passive bandwidth estimation tool, and evaluate their performance in a wireless setting.
- We perform extensive experiments by varying parameters such as interference, rate-adaptation, and number of hops, and identify their impact on the process of estimating Av-Bw in WMNs.
- We evaluate the tools on an indoor testbed, and further validate our results on an outdoor testbed.

The rest of the paper is organized as follows. Section II lists the previous works on estimation of Av-Bw and the motivation behind our work. Section III outlines the details of our experiments, the implementation details of our passive technique and the evaluation methodology. In Section IV, we present the results of the comparison study for both the indoor and the outdoor testbed. Section V concludes the paper.

II. RELATED WORKS & MOTIVATION

Several bandwidth estimation tools have been proposed in past literature. However, not all of these tools focus on measuring Av-Bw in the network. While tools such as Pathchar [4] estimate the per-hop capacity in a network, Nettimer [5] and Pathrate [6] measure the end-to-end capacity. Early bandwidth estimation techniques (such as [7]) suffered from the assumption of fair queuing in routers and as a result cannot estimate the Av-Bw in the current Internet. Tools such as Cprobe [8]

and Pipechar [9] measure a metric called Asymptotic Dispersion Rate (ADR), which is different from Av-Bw [10]. Previous works [11] classified the bandwidth estimation tools into two main categories: the **Probe Gap Model** tools such as Spruce [11] and IGI [12], and the **Probe Rate Model** tools such as Pathload [13], Pathchirp [14] and TOPP [15]. Several previous works have also focused on comparing the performance of existing bandwidth estimation tools [12] [16]. *However, the evaluation in these works has been restricted to wired networks. Our focus is to quantify the performance of these tools in wireless networks.* The theoretical bounds for the capacity of a wireless network [17] are calculated under certain assumptions and do not conform to a real world setting. Several recent works have also proposed using a passive approach for estimating Av-Bw in wireless networks [2] [18] [3]. *However, the performance of these tools has been mostly evaluated via simulations only.* In [19] and [20], the authors have performed limited experiments for 802.11 wireless networks. However, they have not evaluated the impact of factors such as number of hops, packet loss, and rate-adaptation on the bandwidth estimation process.

The lack of understanding of how existing bandwidth estimation tools perform in wireless networks motivates us to perform a thorough comparative study. Some previous works [21] have proposed the use of existing bandwidth estimation techniques in WMNs, without quantifying their performance. We believe that a thorough experimental investigation of the existing tools is of crucial importance in order to validate the accuracy of these tools in real-world wireless networks, and to help guide the design of better monitoring techniques for WMNs.

III. COMPARISON OF EXISTING TOOLS

Before outlining the results from the comparative study, we first describe the tools that we compare, the testbed used, and the validation methodology.

A. Tools & Testbed

We selected four existing probe-based tools: Pathload [13], Spruce [11], PathChirp [14] and PTR [12]. In order to evaluate the relative performance of these tools, we use a testbed in our laboratory, consisting of 802.11a nodes. No other 802.11a networks were detected in the vicinity of our testbed. This gives us the advantage of controlling the amount of cross-traffic and interference in our testbed. The wireless nodes consist of Linux-based laptops and Soekris net4826 embedded devices [22], using Atheros based 802.11 a/b/g wireless cards. We use the open-source wireless driver Madwifi [23] for our testbed. RTS/CTS is disabled for all our experiments. We repeated our experiments for the various physical layer data rates available in 802.11a, since different data rates would result in different channel utilization.

B. Implementation of Passive Tool

Consider a 802.11 network, with RTS/CTS disabled. The following equation shows the total time involved in transmitting a packet:




Fig. 1. 802.11 frame events.

$$T^{CSMA/CA} = T_{DIFS} + T_{backoff} + T_{DATA} + T_{SIFS} + T_{ACK}$$

For a given choice of protocol (802.11 a/b/g), all the values are fixed, except for the back-off time and the data transmission time. The data transmission time depends on the frame size and the modulation rate (physical data rate) used. [24] lists the various parameter values associated with 802.11a scheme. The idea behind the passive scheme is for each node to monitor the channel and listen for ongoing packet transmissions. For each packet that a node hears, it calculates the total transmission time based on the frame size and the data rate used to transmit the packet. Thus, each node can estimate the channel utilization (μ), and multiplying this quantity by the capacity of the channel gives us an estimate of the available bandwidth: $\{BW = (1 - \mu) * C\}$. The capacity (C) of the channel can be estimated by using tools such as Pathrate [20].

In order to obtain this information, we implemented a measurement framework in the Linux kernel, using the madwifi-ng wireless device driver. We modified this device driver to report certain 802.11 events, from which we can derive information about the wireless channel. A user-space program residing on the node, communicates with the modified driver through the NetLink [25] library and records these events. The user-space program can then process these events to obtain the necessary information. Figure 1 shows the temporal relationship between the various 802.11 events reported by the modified driver. The reader can refer to [26] for further details.

C. Evaluation & Validation Methodology

For each of the selected tools, we first ran several experiments to evaluate their performance using the default parameter settings (such as probe size, inter-probe spacing, and number of probes). We wanted to confirm whether the default settings of these tools were also the best for wireless networks. We observed that for some tools the default parameter settings still worked, while for others, choosing a different set of parameters gave more accurate results. All further evaluations are based on the new parameter settings. In order to compare the accuracy of the tools against an accurate benchmark, we leverage the controlled nature of our experiments. Av-Bw is defined as the rate at which a new flow can send data without affecting the rate of the existing flows. In order to estimate Av-Bw, we use the approach outlined in [20]. We term this estimate as the “actual” Av-Bw and use it to estimate the accuracy of the other tools.




Fig. 2. Experimental setup for single hop testing with interference. The solid lines show the actual transmissions, while the dotted lines show that transmissions from node 1 and 3 will interfere at nodes 2 and 4.




Fig. 3. Av. B/W measurement for single-hop with interfering traffic.

IV. EXPERIMENTAL EVALUATION & RESULTS

A. Impact of Interference

For this experiment, we used four laptops (Figure 2), two configured as access points (APs) and the other two as clients. Each of the clients (node 2 and 4) was associated with one of the APs (node 1 and 3 respectively). All the four nodes were in transmission range of each other. No packet loss was observed for this setup at all the transmission rates. The bandwidth estimation tools were run between nodes 1 and 2, while nodes 3 and 4 were used to exchange the cross-traffic (thus creating interference). We fixed the physical data rate of node 3 at 6 Mbps. The cross-traffic was generated at a constant rate of 3 Mbps using a UDP traffic generator. The data rate of the link running the estimation tools was varied from 6 to 54 Mbps (802.11a physical layer data rates).

In Figure 3, for higher modulation rates, it takes less time to transmit the data traffic, and hence the Av-Bw is higher. We can see that Pathload and PTR underestimate the Av-Bw. This is due to 802.11 protocol trying to provide MAC fairness to both the flows. As a result, the packet-train based tools are actually reporting the fair share of the bandwidth, which is lower than the actual Av-Bw. On the other hand, Spruce is over-estimating the bandwidth by large amounts. This happens due to the per-frame sharing of the channel, whereby only a small amount of cross-traffic (typically one frame) gets inserted between the probe pair sent by Spruce. The passive tool returned the most accurate estimate.

Tool	Estimate (Mbps)
Actual	2.34
Pathload	0
Spruce	0
Pathchirp	1.75
PTR	1.625
Passive	2.2

TABLE I
IMPACT OF PACKET LOSS ON ESTIMATION ACCURACY.

B. Impact of Packet Loss

Probe-based estimation tools rely on successful delivery of probe packets from the sender to the receiver. If a significant number of probes are lost, then the tool may give an inaccurate estimate. Hence, we decided to investigate the impact of packet loss on the performance of these tools. In order to have similar rate of packet loss across experiments with different tools, we modified our wireless driver to drop one packet for every fifty correctly received packets. The setup used was the one in Figure 2. The physical layer data rate of both the links was fixed at 6 Mbps. The second pair of nodes (3 and 4) were used to exchange cross-traffic at a constant rate of 3 Mbps. Table I shows that Pathload and Spruce did not report any estimates. This is because both the tools require certain number of probes to get through. For PTR, the loss of probe packets causes it to lower the sending rate, and as a result, it underestimates the Av-Bw. *It should be noted that this scenario also helps us to evaluate the performance of the tools for the case where there is a hidden terminal in the network.* A hidden node would result in packet loss at the receiver. Thus, by simulating the case with packet loss, we are able to quantify the performance of the tools in the presence of a hidden node.

C. Impact of Rate Adaptation

The 802.11 protocol supports multiple physical layer transmission rates depending upon the modulation scheme used. We refer to these as the physical layer data rates or transmission rates. Current 802.11 implementations incorporate a rate-adaptation feature that enables 802.11 radios to vary their data rates based on channel conditions. If the data rate of a node changes, then the channel occupancy time of the node will also change, which will impact the Av-Bw estimation. We evaluate this using the setup shown in Figure 2. The data rate of the link running the bandwidth estimation tools was fixed at 6 Mbps. The second pair of nodes was used to exchange cross-traffic. The data rate of this link is varied over time. We did not use the default rate-adaptation algorithm in 802.11, where the rate changes due to packet loss. This is because it is impossible to repeat the exact experiments using this approach, as we do not get the same variations in the data rate over different runs. In order to model the rate-adaptation behavior, we modified the wireless driver to change the data rate at different instants of time, irrespective of the channel conditions.

Of the selected tools, only Pathchirp and our passive tool provide continuous measurements of Av-Bw. The run time of




Fig. 4. Av. B/W measurement for single hop with rate-adaptation.




Fig. 5. Experimental setup for multi-hop testing. The solid lines show the actual transmissions, while the dotted lines show the self-interference for the same flow between consecutive nodes. The center link was fixed at the lowest data rate of 6 Mbps.

other tools (Pathload, Spruce and PTR) is not user-defined, and they give only one estimate at the end of the measurement. Hence, we could plot only Pathchirp and the passive tool against time. Figure 4 shows the results. The passive tool reacts very quickly to the variations in the data rate. On the other hand, since Pathchirp uses packet trains with varying rates, there is some latency involved before it can adjust its sending rate to the new transmission rate. Pathload estimated the Av-Bw in the range 3.3 - 3.9 Mbps. It seems that since Pathload calculates an average dispersion rate (ADR) before the actual estimation, and the ADR does not get updated when the transmission rate changes, it is unable to reflect the variation in the data rate. PTR gave an average estimate of 2.96 Mbps. Spruce did not give any results and would stop sending packets when the data rate was changed.

D. Impact of Multiple Hops

Multiple wireless hops lead to more issues such as self-interference, which will impact the performance of the selected tools. We decided to investigate this issue with a three hop linear topology (Figure 5). Node 1 and 4 were out of each other’s transmission range. Link 2 interfered with both link 1 and 3. Links 1 and 3 did not interfere with each other. The center link was configured to be on the lowest data rate of 6 Mbps, thus acting as the bottleneck link (Figure 5). The data rates for the other two links were increased from 6 Mbps to 54 Mbps for each run of the experiment. The goal was to see if the various tools are able to identify the bandwidth of the bottleneck link.

We started a UDP flow from node 1 to node 4 at a constant rate of 3 Mbps. We set all the three links to the same 802.11a channel, so that the transmissions on the three hops will




Fig. 6. Av. B/W measurement for three hops with self-interference.

interfere with each other. Figure 6 shows the corresponding results. Multiple hops and self-interference seem to cause large discrepancies in the performance of these tools. The PRM-based tools (Pathload and PTR) once again seem to report the fair share bandwidth, rather than the actual Av-Bw, thus giving lower estimates. Also, as the transmission rate of links 1 and 3 are increased, the channel occupancy time on these links decreases, which increases the bandwidth available on link 2. Spruce does not reflect this increase, and provides erratic results.

E. Evaluation on Outdoor Testbed

From the results of our comparative study, it was clear that a passive technique has several advantages over the currently used probe-based tools. Hence, we decided to further validate the performance of the passive scheme. While our first testbed was setup in a controlled laboratory environment, our second testbed is deployed in a wildlife reserve. It consists of over twenty nodes, spread over two thousand acres of forest land. The testbed uses 802.11a modulation scheme. It has the advantage of being free from any external interference in the region. The traffic in our network comprises of audio and video traffic generated by sensors deployed in the network. This traffic is generated periodically at a fixed rate, and uses fixed sized packets. We did not introduce any additional cross-traffic in the network. This helps us in repeating our experiments with similar background traffic. The reader can refer to [27] for further details on the testbed.

Figure 7(a) shows the results for single-hop. Our passive tool returned fairly accurate estimates. We further tested the performance of the passive scheme with rate-adaptation enabled. For this experiment, we set the data rate of the first link at 6 Mbps, and then changed it to 54 Mbps after a few seconds. Figure 7(b) shows the results estimated by our passive scheme. Our final experiment involved validating the accuracy of the passive scheme for multiple hops. We used a three hop scenario, with each link on the same channel. The center link was fixed at 6 Mbps transmission rate. Figure 7(c) shows that the passive technique provides very close estimates to the “actual” value. These observations conform to the results presented in the previous section, from our laboratory testbed.

V. INFERENCES & CONCLUSIONS

In this paper, we have tackled the problem of estimating Av-Bw in 802.11-based WMNs. Several bandwidth estimation




Fig. 7. Evaluation on OUTDOOR testbed

tools have been proposed previously. Even though these tools were developed for wired networks, several previous works have attempted to use these them on wireless networks. These works lack a thorough evaluation of the performance of the tools on WMNs. Some of the recent works have proposed using a passive approach for bandwidth estimation in wireless networks. However, most previous evaluations have been done via simulations only. We implement our version of this approach by modifying an open source wireless driver, and compare its performance with the previously proposed approaches.

The results from our experiments suggest that probe-based tools are not the best choice for wireless networks. The unique characteristics of wireless networks, such as a shared transmission medium, and the prevailing protocols, such as use of rate-adaptation at the link layer, cause these tools to provide inaccurate results. It was also observed that using these tools over multiple wireless hops causes further degradation in their performance. The passive scheme provided much more accurate results than the probe-based tools.

ACKNOWLEDGMENT

This research was supported in part by the National Science Foundation through the grants CNS-0831914 and CNS-0709264, and by the US Army Research Office through the MURI grant W911NF-07-1-0318.

REFERENCES

- [1] R. Prasad, M. Murray, C. Dovrolis, and K. Claffy, "Bandwidth estimation: metrics, measurement techniques, and tools," *Network, IEEE*, vol. 17, pp. 27–35, 2003.
- [2] H. K. Lee, V. Hall, K. H. Yum, K. Kim, and E. Kim, "Bandwidth estimation in wireless lans for multimedia streaming services," in *IEEE International Conference on Multimedia and Expo (ICME)*, 2006.
- [3] C. Sarr, C. Chaudet, G. Chelius, and I. Lassous, "A node-based available bandwidth evaluation in ieee 802.11 ad hoc networks," in *11th International Conference on Parallel and Distributed Systems (ICPADS)*, 2005.
- [4] V. Jacobson, "Pathchar: a tool to infer characteristics of internet paths," 1997. [Online]. Available: <ftp://ftp.ee.lbl.gov/pathchar>
- [5] K. Lai and M. Baker, "Nettimer: a tool for measuring bottleneck link bandwidth," in *the 3rd USENIX Symposium on Internet Technologies and Systems*, 2001.
- [6] C. Dovrolis, P. Ramanathan, and D. Moore, "Packet-dispersion techniques and a capacity-estimation methodology," *Transactions on Networking, IEEE/ACM*, 2004.
- [7] S. Keshav, "A control-theoretic approach to flow control," in *ACM SIGCOMM Computer Communication Review*, 1995.
- [8] R. Carter and M. Crovella, "Dynamic server selection using bandwidth probing in wide-area networks," Boston University, Tech. Rep., 1996.
- [9] G. Jin, G. Yang, B. Crowley, and D. Agarwal, "Network characterization service (ncs)," in *the 10th IEEE Symposium on High Performance Distributed Computing*, 2001.
- [10] C. Dovrolis, P. Ramanathan, and D. Moore, "What do packet dispersion techniques measure?," in *the 20th IEEE INFOCOM*, 2001.
- [11] J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *the 3rd ACM SIGCOMM IMC*, 2003.
- [12] N. Hu and P. Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," *Selected Areas in Communications, IEEE Journal on*, 2003.
- [13] M. Jain and C. Dovrolis, "Pathload: A measurement tool for end-to-end available bandwidth," in *Passive and Active Measurements workshop (PAM)*, 2002.
- [14] V. J. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell, "Pathchirp: Efficient available bandwidth estimation for network paths," in *Passive and Active Measurements Conference (PAM)*, 2003.
- [15] B. Melander, M. Bjorkman, and P. Gunninberg, "A new end-to-end probing and analysis method for estimating bandwidth bottlenecks," in *IEEE GLOBECOM*, 2000.
- [16] A. Shriram and J. Kaur, "Emperical evaluation of techniques for measuring available bandwidth," in *the 26th IEEE INFOCOM*, 2007.
- [17] P. Gupta and P. R. Kumar, "The capacity of wireless networks," *Information Theory, IEEE Transactions on*, no. 2, Mar 2000.
- [18] S. Shah, K. Chen, and K. Nahrstedt, "Available bandwidth estimation in ieee 802.11-based wireless networks," in *ISMA?CAIDA Workshop on Bandwidth Estimation (BEST)*, 2003.
- [19] M. Neilsen, K. Ovsthus, and L. Landmark, "Field trials of two 802.11 residual bandwidth estimation methods," in *IEEE International Conference on Mobile Adhoc and Sensor Systems (MASS)*, 2006.
- [20] K. Lakshminarayanan, V. Padmanabhan, and J. Padhye, "Bandwidth estimation in broadband access networks," in *the 4th ACM SIGCOMM IMC*, 2004.
- [21] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *MobiCom '04: Proceedings of the 10th annual international conference on Mobile computing and networking*. New York, NY, USA: ACM, 2004, pp. 114–128.
- [22] "Soekris engineering (www.soekris.com)." [Online]. Available: www.soekris.com
- [23] "Multiband atheros driver (www.madwifi.org)." [Online]. Available: <http://madwifi.org>
- [24] "Ieee 802.11a standard." [Online]. Available: <http://standards.ieee.org/getieee802/download/802.11a-1999.pdf>
- [25] C. Benvenuti, *Understanding Linux Network Internals*. O' Reilly, 2006.
- [26] D. Wu, D. P., and P. Mohapatra, "Determining 802.11 link quality with passive measurements," in *IEEE International Symposium on Wireless Communication Systems*, 2008.
- [27] [Online]. Available: <http://spirit.cs.ucdavis.edu/quailridge>