

A Measurement Study of Bandwidth Estimation in IEEE 802.11g Wireless LANs using the DCF

Michael Bredel and Markus Fidler

Multimedia Communications Lab
Technische Universität Darmstadt, Germany
{michael.bredel, markus.fidler}@kom.tu-darmstadt.de

Abstract. In this paper we present results from an extensive measurement study of wireless bandwidth estimation in IEEE 802.11 WLANs using the distributed coordination function. We show that a number of known iterative probing methods, which are based on the assumption of first-come first-serve scheduling, can be expected to report the fair bandwidth share of a new flow rather than the available bandwidth. Our measurement results confirm this view and we conclude that under the current probe gap and probe rate models the fair share can only be loosely related to the available bandwidth. Like a few other studies we report that packet sizes have a tremendous impact on bandwidth estimates. Unlike these studies we can, however, show that minor modifications to known methods for wired networks, such as Pathload, can solve previously indicated limitations of these methods in wireless networks.

1 Introduction

The term available bandwidth denotes the portion of the capacity at a link or a network path that remains unused by present traffic. The idea to estimate the bandwidth of a network path from end-host measurements dates back to TCP congestion control [11] and packet pair probing [17]. Since then the field of available bandwidth estimation has evolved significantly and to date a number of estimation methods exists, e.g. [12, 10, 27, 30, 26], which are frequently used e.g. for network management, error diagnostics, overlay routing, and traffic engineering. The theoretical underpinnings of bandwidth estimation have been explored e.g. in [23, 24, 22, 25] and empirical evaluations can be found e.g. in [29, 30].

The task of bandwidth estimation in wireless networks, such as IEEE 802.11 Wireless LANs, however, has been understood to a much lesser extend. Bandwidth estimation methods, which perform well in case of wired links, have been reported to yield highly unreliable available bandwidth estimates for wireless links [9, 21, 19], hinting at a number of specific challenges and open issues in wireless bandwidth estimation. Here, the performance and the quality of service of a link depend largely on the characteristics of the shared physical medium and the multi-access coordination function. These aspects strongly influence quantities like delay, loss, and throughput and may result in a high variability of

available or actually accessible resources. This makes measurement-based bandwidth estimation in wireless networks a complex and difficult task.

A few approaches to bandwidth estimation specifically address the characteristics of wireless networks. Passive methods, which measure existing traffic, can take advantage of the wireless broadcast medium and record idle periods to estimate the resources that would be available in the proximity of a node [20, 28]. The approach is, however, unreliable in case of hidden stations. Active probing, on the other hand, takes measurements of specific probing traffic at the ingress and the egress of the network to infer the available bandwidth of a network path. Contrary to wired networks a strong impact of packet sizes on bandwidth estimates has been observed for wireless links [15, 16, 21, 19, 6]. Here, the fluid model, which is employed by many estimation methods, is clearly violated. Moreover, it is noted in [19] that the assumption of First-Come First-Serve (FCFS) scheduling, which is the basis of most active probing methods, may not hold in IEEE 802.11 WLANs, e.g. due to the Distributed Coordination Function (DCF).

In this paper we report results from an extensive measurement study of active probing methods to shed light on the issues mentioned above. We conducted measurements of an IEEE 802.11g link in ns-2 simulations [4], Emulab [3, 31], and a highly controlled local wireless testbed that is located in a shielded and reflection absorbing measuring room. Based on our measurement results we find that the common FCFS model does clearly not hold under the DCF. The DCF can (within certain limits) rather be viewed as implementing fair scheduling. Based on the fair queuing model we can show that iterative methods that are based on the FCFS assumption can be expected to estimate the fair share of a new flow instead of the available bandwidth. Our measurements using the methods in Tab. 1 support the anticipated results. The gathered data reconfirms the known dependency of bandwidth estimates on the probing packet size.

The remainder of this paper is organized as follows. In Sect. 2 we review the state-of-the-art in available bandwidth estimation. In Sect. 3 we discuss relevant characteristics of wireless links and the fair queuing model. In Sect. 4 we show our measurement results and in Sect. 5 we present our conclusions.

2 Methods for Available Bandwidth Estimation

In this section we discuss the state-of-the-art of available bandwidth estimation in wired and wireless networks. We focus on a set of publicly available measurement tools, see Tab. 1, which are used for measurements in Sect. 4. For related empirical evaluations in wired networks see e.g. [29, 30].

The task of available bandwidth estimation is to infer the portion of the capacity of a link or a network path that remains unused by cross-traffic. The available bandwidth of a link with index i can be defined as [12]

$$AB_i(\tau, t) = C_i(1 - u_i(\tau, t)) \quad (1)$$

where C_i is the capacity and $u_i \in [0, 1]$ is the utilization by cross-traffic in the interval $[\tau, t)$. The available bandwidth of a network path is determined by the available bandwidth of the tight link as $AB(\tau, t) = \min_i\{AB_i(\tau, t)\}$ [12].

Table 1. Bandwidth estimation methods used in this study

Probing method	Probing traffic	Inference technique
Pathload [12]	packet trains	iterative
Pathchirp [27]	packet chirps	iterative
Spruce [30]	packet pairs	direct
IGI [10]	packet trains	direct
PTR [10]	packet trains	iterative
WBest [21]	packet trains	direct
DietTOPP [14]	packet trains	iterative

Active measurement methods inject specific probes into the network and estimate the available bandwidth from measurements of the probing traffic at the ingress and at the egress of the network. The majority of the methods uses packet pairs, i.e. two packets sent with a defined spacing in time referred to as gap, or packet trains, i.e. a larger number of packets sent at a defined constant rate. The rate of a packet train can be converted into a certain spacing of the train’s packets, showing a direct relation to the gap model of packet pairs. Packet chirps [27] are specific packet trains that are sent at a geometrically increasing rate respectively with a geometrically decreasing gap.

Many methods use a simplified network model, where cross-traffic is viewed as constant rate fluid and the network is abstracted as a single tight link. Under these assumptions the available bandwidth of a network path simplifies to $AB = C(1 - u)$. In addition, FCFS multiplexing is usually assumed, where flows share the capacity of a link proportionally to their offered rates. For constant rate probes an expression referred to as rate response curve [26, 23] can be derived as

$$\frac{r_i}{r_o} = \max\left(1, \frac{r_i + \lambda}{C}\right) = \begin{cases} 1 & , \text{ if } r_i \leq C - \lambda \\ \frac{r_i + \lambda}{C} & , \text{ if } r_i > C - \lambda \end{cases} \quad (2)$$

where r_i and r_o are the input and output rates of probes respectively and λ is the input rate of cross-traffic. If $\lambda \leq C$ the available bandwidth follows as $AB = C - \lambda$ and otherwise $AB = 0$. Based on this model the task of available bandwidth estimation is to select the rate of probing traffic such that (2) can be solved for C and λ or $C - \lambda$. While (2) is usually used for packet train probes, an equivalent gap response curve can be derived for packet pairs, where the gap g is linked to the rate r by the packet size l resulting in $g_i = l/r_i$ and $g_o = l/r_o$ [23].

In [13] measurement methods are classified by their inference technique as either *direct* or *iterative* probing schemes. Direct probing schemes assume that the capacity of the link C is known in advance. In this case (2) can be solved for the rate of the cross-traffic if the probing rate is larger than the available bandwidth. A straightforward choice is to probe with $r_i = C$ in which case the available bandwidth follows from (2) in rate respectively gap notion as [21, 30]

$$AB = C \left(2 - \frac{C}{r_o}\right) = C \left(1 - \frac{g_o - g_i}{g_i}\right). \quad (3)$$

Spruce, WBest, and IGI are methods that use direct probing. Spruce assumes that the capacity is known a priori and immediately applies the gap version of (3). WBest provides a two-step algorithm using packet pairs to estimate the link capacity and packet trains for available bandwidth estimation based on the rate version of (3). IGI uses probing trains with increasing gaps resulting in a more complex direct probing formula than (3), for details see [10].

Iterative probing methods do not require a priori knowledge of the link capacity. They employ an iterative procedure with multiple probing rates aiming to locate the turning point of the rate response curve (2), i.e. they seek to find the largest probing rate r_i such that $r_i/r_o = 1$. At this point the probing rate coincides with the available bandwidth.

TOPP, DietTOPP, PTR, Pathload, and Pathchirp are iterative probing methods. TOPP [26] uses trains of packet pairs with increasing rate and applies (2) for available bandwidth estimation. It recursively extends the model to the multiple node case and in addition it estimates the capacity from the second linear segment in (2). Closely related is a simplified version called dietTOPP. PTR is a packet train method that uses a gap version of (2).

Pathload varies the rate of packet trains using a binary search algorithm to find the largest probing rate that does not cause overload and hence matches the available bandwidth. It uses increasing one-way delays as an indication of overload. Increasing delays indicate that the input rate exceeds the output rate, i.e. $r_i/r_o > 1$ which clearly shows the relation to (2) [23]. Pathchirp increases the probing rate within a single packet train, referred to as a chirp, instead of varying the rate of successive packet trains. Like Pathload it detects crossing the turning point of the rate response curve from increasing one-way delays.

Most of the discussed methods have been developed for wired networks while WBest and dietTOPP have been suggested by the authors for available bandwidth estimation in wireless networks. Further on, a method called ProbeGap has been proposed for bandwidth estimation in broadband access networks [19]. The method does not exactly fit into the classification scheme used here. ProbeGap sends out single packets and collects the one-way delays of these probes. The fraction of the packets which have a delay close to zero are assumed to have found an idle channel. This fraction is used to estimate the available bandwidth. Besides, passive measurement approaches can take advantage of the wireless broadcast medium [28] or protocol related information [20]. Passive methods are, however, not considered within the scope of this study.

3 Relevant Wireless Link Characteristics

In this section we discuss relevant characteristics of wireless links that are of vital importance for bandwidth estimation. We show how these aspects affect current fluid rate and gap models from Sect. 2 and reason which quantity we expect to be estimated by known methods for bandwidth estimation in wireless systems. We use the term wireless link meaning a wireless broadcast channel with Medium Access Control (MAC) and Radio Link Control (RLC) protocols.

Fading and interference: As opposed to wired links the characteristics of wireless channels are highly variable due to fading. Other potentially hidden stations, which may even include stations that implement different radio standards using the same frequency band, create interference on the wireless broadcast medium. These effects can cause rapid fluctuations of the signal-to-noise ratio and may lead to high bit error rates. Different modulation and coding schemes combined with rate adaptation may be used for compensation. As a consequence, the capacity and the availability of the channel may vary drastically.

Contention: In case of wireless multi-access channels, stations share the same medium and contend for access to the channel, often in a fully distributed manner. Channel access is controlled by the MAC protocol. Before accessing the medium stations listen to the channel to detect nearby transmissions with the objective of avoiding collisions. This procedure may fail in case of hidden stations, thus requiring additional protocol mechanisms such as RTS/CTS. The resulting behavior of medium access procedures may be largely different if compared to FCFS multiplexing at a point-to-point link.

Retransmissions: Due to frequent packet loss on wireless links, e.g. because of fading, interference, or collisions, many standards include an RLC protocol that implements an automatic repeat request technique such as stop-and-wait ARQ to ensure packet delivery. Link layer retransmissions consume channel capacity and lead to increased and varying one-way delays.

Effects that are due to fading and interference result in a time-varying channel capacity $C(t)$. It is straight-forward to adapt the definition of available bandwidth (1) accordingly. The fluid rate and gap models for bandwidth estimation (2), (3) assume, however, a constant capacity. This assumption is mirrored by the view of cross-traffic as constant rate fluid.

Our controlled measurement environment eliminates effects that are due to fading and interference to a large extend. Protocol-related aspects, however, remain to be addressed. In the following subsections we investigate two effects of the IEEE 802.11 protocols that have major impact on bandwidth estimation.

3.1 Protocol Overhead

We compute the impact of the IEEE 802.11g protocol overhead on the achievable throughput. Given packets of size l bits the throughput can be determined as $C = l/g$ where g is the gap, i.e. the time between the beginning of two subsequent packets transmitted at the maximum rate. The gap consists of a number of additive parameters¹ including inter frame spacings, the expected backoff time, transmission times of the data packet including preamble and header and the acknowledgement, as well as propagation delays which are neglected here.

Fig. 2 shows the achievable throughput for different packet sizes. The outcome of the above model is compared to measurement results from our testbed shown in Fig. 1, for details see Sect. 4. We used a single greedy UDP traffic stream that is generated with the D-ITG traffic generator [1]. The measured throughput is

¹ We set $t_{\text{difs}} = 28 \mu\text{s}$, $t_{\text{sifs}} = 16 \mu\text{s}$, $t_{\text{backoff}} = 139.5 \mu\text{s}$, $t_{\text{data}} = 26 + l/54 \mu\text{s}$, $t_{\text{ack}} = 30 \mu\text{s}$




Fig. 1. Wireless testbed setup.




Fig. 2. Achievable throughput.

averaged over 60 s. It exhibits a strong dependence on the packet size, as also derived for IEEE 802.11b in [8].

We conclude that the size of probing packets used by bandwidth estimation tools has a large impact on the accuracy of estimates. This aspect is also noted in [16, 19, 21, 6]. The fluid rate response model for bandwidth estimation does not involve a notion of packets. Using the equivalent gap formulation, an extended model that includes the effects of the packet size has been developed in [16]. The conclusion in [16] is none the less that the probing packet size needs to be tailored to the application that uses the bandwidth estimates.

3.2 Distributed Coordination Function

In this section we investigate how the capacity of IEEE 802.11g links using the DCF is allocated to contending flows. The issue is discussed controversially in the literature. While [18] showed short-term unfairness of CSMA/CA-based WaveLANs a recent study [7] attributes findings of unfairness to early WaveLAN cards and reports that current IEEE 802.11 DCF implementations actually exhibit good short-term fairness. On this account we performed a number of initial experiments to explore fairness issues in our shielded testbed.

Fig. 3 shows the throughput of contending flows at an IEEE 802.11g link. Flows have a constant bit rate and are generated using the Rude/Crude traffic generator [5]. All packets have a constant size of 1500 Bytes. The throughput is averaged over 60 s. In Fig. 3(a) two flows contend for the link. Flow 1 has a rate of 28 Mbps and the rate of flow 2 is increased from 0 to 28 Mbps in steps of 1 Mbps after each experiment. Similarly Fig. 3(b) shows the throughput of four flows, where the rate of flow 4 is increased.

The results in Fig. 3 confirm that each flow receives a fair share of the capacity. Flow 2 in Fig. 3(a) achieves its target throughput whereas the throughput of flow 1 is reduced accordingly until flow 2 reaches 14 Mbps. From this point on both flows get a fair share of 14 Mbps regardless of the rate of flow 2. Fig. 3(b) confirms this result for four heterogeneous flows. Note that we measured fair throughput shares only for flows with homogeneous packet sizes.

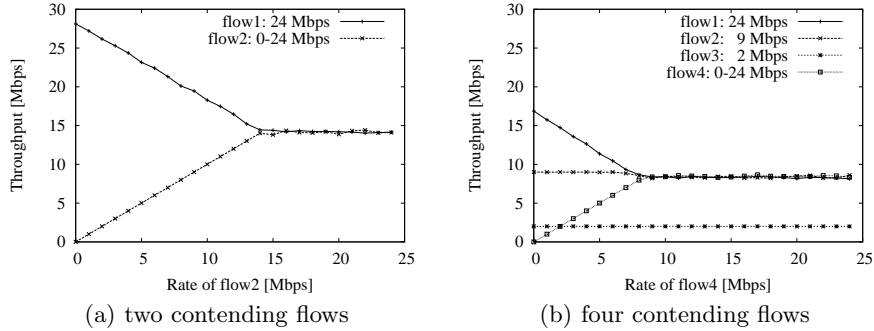


Fig. 3. Measured throughput of contending flows. Each flow obtains its fair share.

The fair share f at a congested link that is observed in Fig. 3 can be computed as the solution of

$$f : \sum_{k=1}^n \min\{r_k, f\} = C \quad (4)$$

where C is the capacity, r_k is the rate of flow k , and n is the number of flows. Once f is determined the output rate of flow k follows as $\min\{r_k, f\}$.

The rate response curve of a fair queuing system follows immediately as

$$\frac{r_i}{r_o} = \max\left(1, \frac{r_i}{f}\right) = \begin{cases} 1 & , \text{ if } r_i \leq f \\ \frac{r_i}{f} & , \text{ if } r_i > f \end{cases} \quad (5)$$

where r_i and r_o are the input and output rates of the probes respectively. As opposed to the FCFS rate response curve (2) the available bandwidth cannot be derived from (5). Trivially the available bandwidth AB is upper bounded by the fair share, i.e. $0 \leq AB \leq f$. Without further assumptions the two extremal values can, however, be easily attained if $f \leq C/2$. As an example consider a single contending flow with rate $\lambda = C/2$ respectively $\lambda = C$. The fair share of a new greedy flow is $f = C/2$ in both cases, whereas the available bandwidth becomes $AB = f$ respectively $AB = 0$.

Referring to the classification of bandwidth estimation methods in Sect. 2 we conclude that iterative methods, which use the turning point of the rate response curve as bandwidth estimate, can be expected to report the fair share of a new greedy flow in case of a fair wireless link. For existing direct probing methods that inject probes with rate $r_i = C$ such a clear result cannot be established. Inserting $r_o = f$ into (3) does neither compute the available bandwidth nor the fair share. We note, however, that direct probing with $r_i = C$ could easily report the fair share, since $r_o = f$ in this case.

4 Experimental Evaluation of Bandwidth Estimation

Equipped with the results from Sect. 3 we now investigate the performance of the bandwidth estimation tools listed in Tab. 1. If not mentioned otherwise we use the default configuration of the bandwidth estimation tools to perform the experiments. We evaluate the methods using a wireless testbed in a shielded, anechoic room. Hence, we act on the assumption that the physical medium is free of interference from external sources that do not belong to the testbed. We focus on the accuracy of wireless bandwidth estimates and show how these relate to the available bandwidth respectively to the fair share under different types of contending traffic. We do not report probing overhead, intrusiveness, as well as run or convergence times. These aspects are elaborated e.g. in [29].

As shown in Fig. 1 the testbed comprises five wireless stations (S1 to S5) that serve as traffic sources. The stations are connected to the access point (AP) using IEEE 802.11g with 54 Mbps². The access point is connected over fast Ethernet at 100 Mbps to a station (R) that acts as receiver. The distance between the wireless stations and the access point was between 0.5 m and 1.5 m. We switched off RTS/CTS, automatic rate adaption (which turned out to be unneeded in our scenario) as well as packet fragmentation. We used the DCF for medium access. In parallel to the wireless test network, all nodes are connected to a separated switched Ethernet, which is used as a control network.

4.1 Impact of the Intensity of Contending Traffic

In the first set of experiments we estimate the available bandwidth from S1 to R in the presence of a single contending flow. We increase the rate λ of the contending traffic that flows from S2 to R from 0 Mbps up to 28 Mbps in steps of 1 Mbps. The contending traffic consists of packets of 1500 Bytes and is generated using the D-ITG traffic generator [1]. All probe packets are set to 1500 Bytes.

Fig. 4 shows the average of 25 available bandwidth estimates for each of the tools and all rates of the contending traffic as well as corresponding confidence intervals at a confidence level of 0.95. As a reference the available bandwidth $AB = C - \lambda$ as well as a the fair share of a new flow $f = \max\{C - \lambda, C/2\}$ are plotted, where we use $C = 28$ Mbps from Fig. 2 for a packet size of 1500 Bytes.

Iterative probing: From our arguments in Sect. 3 we expect that the iterative probing methods Pathload, DietTOPP, Pathchirp, and PTR report an estimate of the fair share. As indicated in Fig. 4 the fair share and the available bandwidth are identical for contending traffic with rate $\lambda \in [0 \dots 14]$ Mbps, whereas they differ for $\lambda \in (14 \dots 28]$ Mbps. Fig. 4(a) shows that the estimates from Pathload (which reports an upper and a lower bound of the available bandwidth) and DietTOPP clearly confirm the fair queuing model in (5). Both methods closely track the fair share and the reported estimates deviate noticeably

² We used Lenovo ThinkPad R61 notebooks with 1.6 GHz, 2 GB RAM running Ubuntu Linux 7.10 with kernel version 2.6.22. We employed the internal Intel PRO/Wireless 4965 AG IEEE 802.11g WLAN adapters. The access point is a Buffalo Wireless-G 125 series running DD-WRT [2] version 24 RC-4.

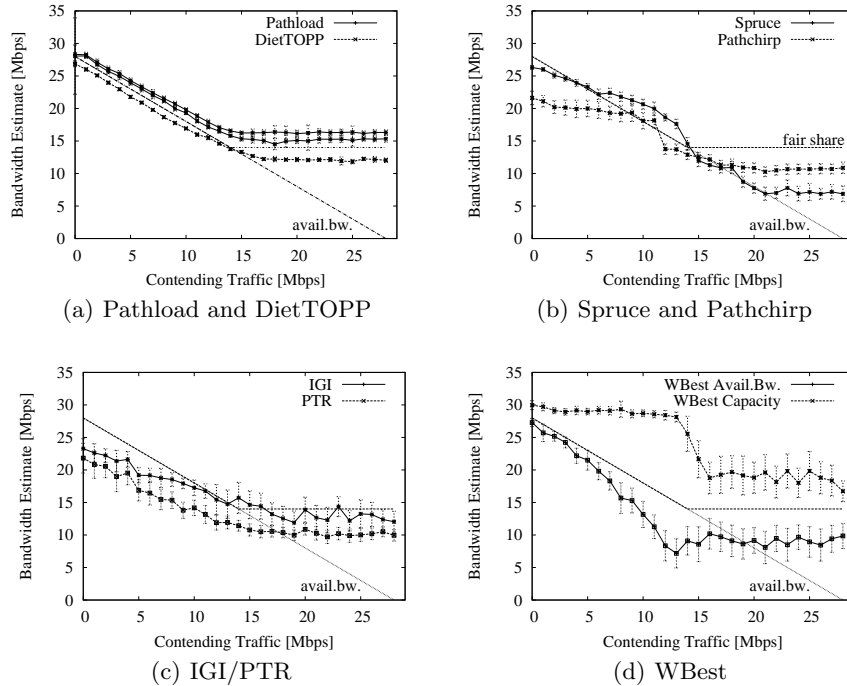


Fig. 4. Bandwidth estimates for a wireless link with one contending flow.

from the available bandwidth as the rate of the contending traffic increases beyond 14 Mbps. The results from PTR in Fig 4(c) and to a lesser extent from Pathchirp in Fig. 4(b) confirm this view. In case of Pathchirp we used the estimates provided after Pathchirp’s self-adapting phase. In our experimental results plotted in Fig. 4(b) it nevertheless underestimates the fair share and the estimates exhibit a comparably high variance.

Pathload has been reported to provide inaccurate bandwidth estimates for wireless networks in [9, 19, 21]. This stands in contrast to experiences made using Pathload in wired networks. In [21] the probing packet size is mentioned as a possible reason for bandwidth underestimation, and [19] identifies the signature of one-way delays in wireless networks as a source of the problem. Having confirmed the strong impact of the packet size on the throughput in wireless networks, see Fig. 2, we modified Pathload so that we can specify the probing packet size. Using Pathload with a fixed packet size of 1500 Bytes improves bandwidth estimates significantly and leads to quite accurate and stable results as can be seen in Fig. 4(a). Similar problems have been reported for IGI/PTR [21], which can also be mitigated using a packet size of 1500 Bytes rather than the default size. The estimates in Fig. 4(c) are, however, less sensitive to the intensity of contending traffic, as also reported for cross-traffic in wired networks in [30].

Direct probing: Direct probing tools require a priori knowledge of the link capacity. We executed Spruce with a given capacity of $C = 28$ Mbps which corresponds to the throughput for packets of 1500 Bytes size in Fig. 2. From our results shown in Fig. 4(b) we cannot detect a clear trend of the estimates towards either the fair share or the available bandwidth once λ exceeds 14 Mbps.

WBest uses a two-step algorithm to estimate first the capacity and then the available bandwidth. In our measurements both estimates exhibit a comparably high variance as shown in Fig. 4(d). Moreover, the capacity estimates are sensitive to contending traffic, possibly a result of fair resource allocation, such that bandwidth estimates that are based heron may be unreliable.

4.2 Impact of the Number of Contending Flows

In the second set of experiments we investigate the impact of the number of contending flows on bandwidth estimates. We use contending traffic with a total rate of 20 Mbps, which is divided evenly among one to four flows. Hence, the available bandwidth $AB = 8$ Mbps remains constant in all experiments, whereas the fair share of a new flow is $(14, 9.\bar{3}, 8, 8)$ Mbps for $(1, 2, 3, 4)$ contending flows each offering a rate of $(20, 10, 6.\bar{6}, 5)$ Mbps respectively. Since we only related the estimates of iterative probing methods to the fair share we restrict the results shown here to iterative methods. Again all contending and probe packets are adjusted to have a fixed size of 1500 Bytes to achieve comparability.

As in the experiments presented above, we ran each experiment 25 times and display the average of the available bandwidth estimates and belonging confidence intervals at a confidence level of 0.95 in Fig. 5. The estimates of the iterative methods Pathload, DietTOPP, PTR, and Pathchirp are closely related to the fair share and do not match the available bandwidth. As anticipated these results confirm the model that is developed in Sect. 3.

Without presenting results due to limited space we state that the direct probing tools that were investigated tend to be more inaccurate and do not report the available bandwidth nor the fair share in these experiments.

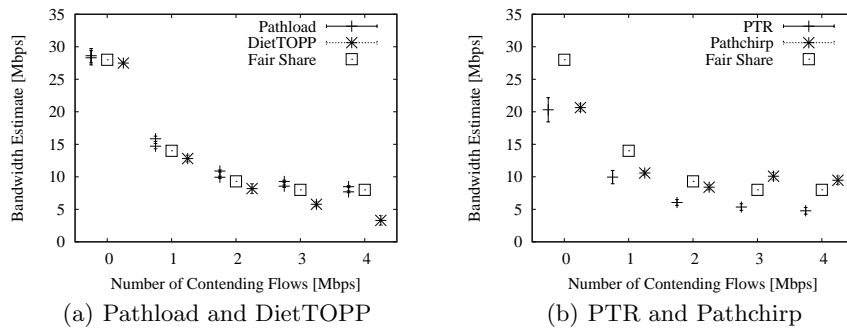


Fig. 5. Bandwidth estimates for a wireless link with several contending flows.

5 Conclusions

We conducted an extensive measurement study of wireless bandwidth estimation in IEEE 802.11g WLAN testbeds. In contrast to wired links bandwidth estimates for wireless channels depend largely on the choice of packet sizes. We adapted the examined tools accordingly. We found that the FCFS assumption common in bandwidth estimation does not apply in case of wireless channels with contending traffic, where the distributed coordination function seeks to achieve a fair bandwidth allocation. We showed that the estimates of known iterative measurement methods can be related to the fair share of a new flow, which may deviate significantly from the available bandwidth. Our measurement results confirm this relation. A similar result was not established for direct probing. As opposed to previous studies our measurement results indicate that the methods, which have been specifically targeted at wireless channels, are not superior to previously known methods that have been developed for wired networks.

Acknowledgements

We would like to thank M. Hollick, R. Steinmetz, and R. Jakoby for providing the equipment and facilities that made this research possible. This work has been funded by an Emmy Noether grant of the German Research Foundation.

References

1. D-ITG: Distributed Internet Traffic Generator. www.grid.unina.it/software/ITG.
2. dd-wrt: Linux wireless router. www.dd-wrt.com.
3. Emulab: Network emulation testbed. www.emulab.org.
4. ns-2: Network simulator. www.isi.edu/nsnam/ns.
5. Rude/Crude: Real-time UDP Data Emitter/Collector. rude.sourceforge.net.
6. A. Amambra, K. M. Hou, and J.-P. Chanet. Evaluation of the performance of the SLoPS: Available bandwidth estimation technique in IEEE 802.11b wireless networks. In *Proc. of IFIP NTMS*, pages 123–132, May 2007.
7. G. Berger-Sabbatel, A. Duda, M. Heusse, and F. Rousseau. Short-term fairness of 802.11 networks with several hosts. In *Proc. of IFIP MWCN*, pages 263–274, Oct. 2004.
8. G. Bianchi. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE JSAC*, 18(3):535–547, Mar. 2000.
9. A. Botta, A. Pescapé, and G. Ventre. Improving accuracy in available bandwidth estimation for IEEE 802.11-based ad hoc networks. In *Proc. of IEEE Systems Communications*, pages 287–292, Aug. 2005.
10. N. Hu and P. Steenkiste. Evaluation and characterization of available bandwidth probing techniques. *IEEE JSAC*, 21(6):879–894, Aug. 2003.
11. V. Jacobson. Congestion avoidance and control. In *Proc. of ACM SIGCOMM*, pages 273–288, Aug. 1988.
12. M. Jain and C. Dovrolis. End-to-end available bandwidth: measurement methodology, dynamics, and relation with tcp throughput. *IEEE/ACM TON*, 11(4):537–549, Aug. 2003.

13. M. Jain and C. Dovrolis. Ten fallacies and pitfalls on end-to-end available bandwidth estimation. In *Proc. of ACM IMC*, pages 272–277, Oct. 2004.
14. A. Johnsson, B. Melander, and M. Björkman. Diettopp: A first implementation and evaluation of a simplified bandwidth measurement method. In *Proc. of SNCNW*, Nov. 2004.
15. A. Johnsson, B. Melander, and M. Björkman. Bandwidth measurement in wireless networks. In *Proc. of Med-Hoc-Net*, June 2005.
16. A. Johnsson, B. Melander, and M. Björkman. An analysis of active end-to-end bandwidth measurements in wireless networks. In *Proc. of IEEE/IFIP E2EMON*, pages 74–81, Apr. 2006.
17. S. Keshav. A control-theoretic approach to flow control. In *Proc. of ACM SIGCOMM*, pages 3–15, Sept. 1991.
18. C. E. Koksal, H. Kassab, and H. Balakrishnan. An analysis of short-term fairness in wireless media access protocols. In *Proc. of ACM SIGMETRICS*, pages 118–119, June 2000.
19. K. Lakshminarayanan, V. N. Padmanabhan, and J. Padhye. Bandwidth estimation in broadband access networks. In *Proc. of ACM IMC*, pages 314–321, Oct. 2004.
20. H. K. Lee, V. Hall, K. H. Yum, K. I. Kim, and E. J. Kim. Bandwidth estimation in wireless LANs for multimedia streaming services. In *Proc. of IEEE ICME*, pages 1181–1184, July 2006.
21. M. Li, M. Claypool, and R. Kinicki. WBest: A bandwidth estimation tool for multimedia streaming application over IEEE 802.11 wireless networks. Technical Report WPI-CS-TR-06-14, Computer Science Department, Worcester Polytechnic Institute, Mar. 2006.
22. J. Liebeherr, M. Fidler, and S. Valaee. A min-plus system interpretation of bandwidth estimation. In *Proc. of IEEE INFOCOM*, pages 1127–1135, May 2007.
23. X. Liu, K. Ravindran, and D. Loguinov. A queuing-theoretic foundation of available bandwidth estimation: Single-hop analysis. *IEEE/ACM Trans. Networking*, 15(4):918–931, Aug. 2007.
24. X. Liu, K. Ravindran, and D. Loguinov. A stochastic foundation of available bandwidth estimation: Multi-hop analysis. *IEEE/ACM Trans. Networking*, 16(2), Apr. 2008.
25. S. Machiraju, D. Veitch, F. Baccelli, and J. Bolot. Adding definition to active probing. *ACM SIGCOMM Computer Communication Review*, 37(2):19–28, Apr. 2007.
26. B. Melander, M. Björkman, and P. Gunningberg. Regression-based available bandwidth measurements. In *Proc. of SPECTS*, July 2002.
27. V. J. Ribeiro, R. H. Riedi, R. G. Baraniuk, J. Navratil, and L. Cottrell. Pathchirp: Efficient available bandwidth estimation for network paths. In *Proc. of PAM*, Apr. 2003.
28. C. Sarr, C. Chaudet, G. Chelius, and I. G. Lassous. Improving accuracy in available bandwidth estimation for IEEE 802.11-based ad hoc networks. In *Proc. of IEEE MASS*, pages 517–520, Oct. 2006.
29. A. Shriram and J. Kaur. Empirical evaluation of techniques for measuring available bandwidth. In *INFOCOM 2007. 26th IEEE International Conference on Computer Communications.*, pages 2162–2170. IEEE, May 2007.
30. J. Strauss, D. Katabi, and F. Kaashoek. A measurement study of available bandwidth estimation tools. In *Proc. of ACM IMC*, pages 39–41, Oct. 2003.
31. B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. In *Proc. of USENIX OSDI*, pages 255–270, Dec. 2002.