Systematic Measurement of TCP Performance over Wireless LANs

B. Rathke, M. Schläger, A. Wolisz

Technical Report TKN-01BR98

Technical University Berlin, Sekr. FT 5-2, Einsteinufer 25, 10587 Berlin, Germany
{rathke, morten, wolisz}@ee.tu-berlin.de

Abstract

One challenge in the development of telecommunication networks is the seamless integration of wireless devices into the global Internet. Although it is well known that a wireless end-system suffers from the characteristics of the wireless link systematic measurements are still missing. In this paper we present results of a systematic measurement of TCP performance over wireless LANs. In this measurement we investigated location dependent throughput of a whole wireless radio cell. We used pre 802.11 compatible wireless LAN technologies with link layer error control and without.

1. Introduction

It is well known that TCP suffers from the loss and delay characteristics of wireless links. Several approaches have been presented to improve the performance. However most of the investigation were based on simulation or measurement with artificial assumptions about the wireless link (e.g error generator) [1]. In contrast to this the results presented in this technical report are obtained by measurements in a real environment without any additional error generator. Such measurements are not only important for understanding TCP’s behavior over WLANs but are needed for more accurate and realistic simulations.

We used a wireless LAN which offers a unreliable service as well as a wireless LAN which offers a reliable service. A reliable service means that the MAC of the wireless LAN includes some kind of error control mechanisms (e.g. immediate acknowledgment). In unreliable wireless LANs the error recovery task is up to the higher layers. Although the current trend is to use reliable MACs, as can be seen from the fact that both standards for wireless LANs (IEEE 802.11 [4] and Hiperlan [3]) define a reliable MAC, there are also arguments that speak for the use of an unreliable MAC. The motivation for using a reliable MAC is the error prone characteristic of wireless links. An efficient error control mechanism like immediate acknowledgment might be important for performance considerations. On the other hand it has been observed recently that the error characteristics of WLANs are bursty [8]. This means that if a packet is corrupted the probability that the next several packets are corrupted as well is high. Under this assumption the immediate acknowledgment wastes scarce resources instead of offering a high performance. Because of these two different points of view it can not be foreseen which technology will be favored. Therefore we considered both technologies in our investigations.

We use ARLAN [10] as a wireless LAN with reliable MAC service and WaveLAN [9] as a wireless LAN with unreliable MAC service. Both WLANs have in common that they provide a bit rate of 2 Mbit/s, that they use CSMA/CA as medium access mechanism and that they are based on DSSS operating at 2,4 GHz. The features of ARLAN are quite similar to the features of 802.11. Therefore performance results obtained by using ARLAN are a first indication for the operation of 802.11 compatible WLANs. A performance study of 802.11 and Hiperlan based on simulations can be found in [6].

The article is structured as follows. In the second chapter we introduce our measurement set-up. In third chapter we present and discuss the obtained results. The paper is completed by final remarks.

2. Experimental Environment

Our configuration for the performance evaluation is depicted in figure 1. This structure is used for all performed measurements. It consists of a wireless end-system which communicates via a bridge (base-station) with a fixed host. The bridge and the fixed end-system are connected with a 10 Mbit/s Ethernet. The wireless LAN technology used for the connection between the wireless host and the bridge is either WaveLAN or ARLAN. No other stations are connected to either the fixed or the wireless part of the network and there is no overlapping with other WLAN radio cells. To obtain reproducible measurements results, we eliminated disturbing effects from other end-systems either from the wired or wireless network.

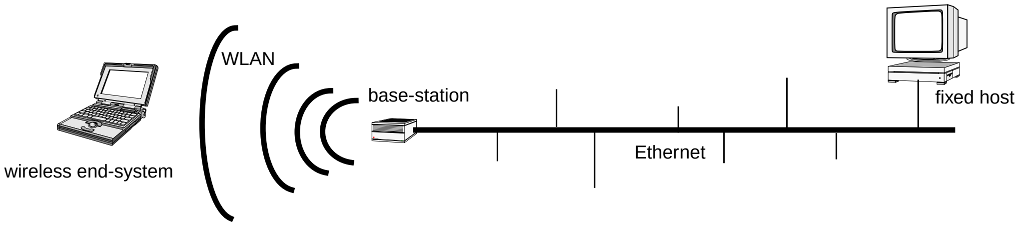

Figure 1:General measurement configuration

FTP was used as the application. The FTP client, running on the wireless end-system requests files from the FTP server running on the fixed host. The size of the requested file was either 570 kbytes to investigate long time behavior or 30 kbytes to show short time characteristics. We execute a series of file transmissions called a trial. In case of long files each trial consist of 10 file transfers and in case of small files each trial consists of 100 file transfers. We measured the throughput on application level for each transfer. Additionally we measured the loss rate. Beside these results we observed a lot more variables of the TCP context record at the fixed host (e.g. congestion window, advertised window, estimated round trip time). For this purpose we developed a tool called SNUFFLE [5]. The variables give us the opportunity for an appropriate interpretation of the achieved throughput by each architecture.

The end-systems used are PC’s with an Intel Pentium processor with 133 MHz running under Linux version 2.0.27. The Linux TCP implementation is similar to the Reno TCP implementations of BSD. However the Linux TCP differs from the BSD implementation in two cases. Linux TCP counts the congestion window in multiple of segments whereas BSD counts the congestion window in bytes. Further the retransmission timer granularity is smaller (10ms) compared to BSD (500ms) and the smallest value of the Linux retransmission timer is 200ms. BSD uses a minimum value of 500ms.

The base-station was developed by us and is also implemented on a PC with an Intel Pentium processor (133MHz) running under Linux 2.0.27. Its functionality at the current state is simple. The base-station can support different modes. It is transparent to the end-system (bridging functionality) or integrates the Snoop[1] agent. Currently the base-station does not support handovers. Using the transparent bridging mode, our base-station achieves performance results similar to commercial available base-stations.

The place of investigation is a single floor of the Electrical Engineering Department where our labs and terminal rooms are located (see figure 2). The floor and the whole building was built in the sixties, in the typical way, which has some features making the radio transmission fairly difficult:

---

[1] Results obtained using Snoop will follow in a upcoming paper.

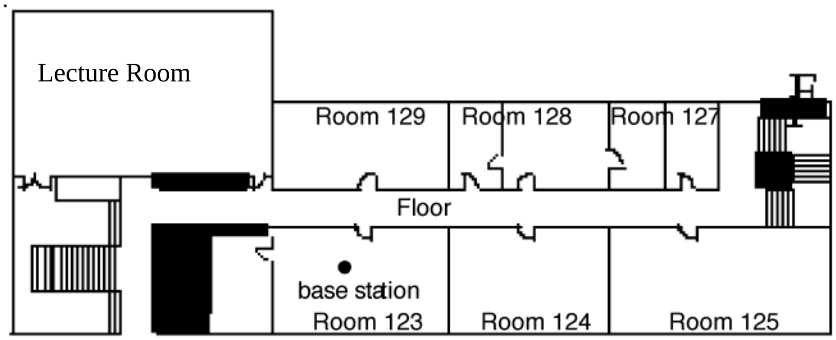

Figure 2:Measurement environment

- The walls are out of ferro-concrete.
- In the floor are cupboards with steel doors.
- Along the walls with doors as well as along the window walls are data cabling systems and electrical power cabling systems.
- The rooms are equipped with a lot of workstations.

We admit that those conditions are not easy, but we argue that they are not uncommon. Preliminary results demonstrated that the whole floor can be considered as a single radio cell covering an area around 40x50 square meters.

The base-station was positioned in the middle of room 123. During each trial, as described above, the wireless host resides at a fixed location within the radio cell. After the trial is completed, the wireless host will be positioned at another location and the next trial starts.

3. Measurements Results

The measurements are done in two steps. First, we performed an systematic measurement of a whole floor at our building using TCP. Due to the location dependent characteristics of WLANs, we are interested in the operation of TCP at different positions. We are not looking for the physical reasons of the performance but concentrate on effects seen by TCP. In a second step we selected typical positions from the initial measurements to obtain a deeper understanding of the TCP behavior.

3.1. Location dependent throughput characteristics

The radio cell has been treated as a grid with a resolution of 2 meters, and the trials have been repeated systematically in each node of this grid. As expected the mean TCP-throughput depends strongly upon the location, and varies from zero up to around 120 kbytes/s for ARLAN and up to 180 kbytes/s for WaveLAN. Further we observed that the dependency is similar for both technologies (ARLAN and WaveLAN). They only differ in the maximum achievable throughput. Therefore, we describe only the results obtained by using ARLAN.

The highest throughput was achieved in rooms 123 to 125. On the floor and in rooms 127 to 129 the throughput is substantially reduced, in room 127 there are even areas, where no data transmission is possible. It is even wrong to expect a homogeneous throughput on each side of the corridor, in rooms 123-125, even in room 123 itself the differences in throughput are non-negligible and definitely not regular.

For the sake of analysis of this phenomena, we classified the locations into three types, those having mean throughput higher than 100 kbytes/s, in the range of 30-100 kbytes/s and below 30 kbytes/s and called a good position, a bad position and an unacceptable position, respectively. For the in depth analysis we have selected several positions of each type. For each position we have investigated essential questions. On one hand, we want to show, whether the decrease in throughput is caused by the MAC (retransmissions of ARLAN / losses of WaveLAN) or rather caused by the TCP's mechanisms. An upper bound for TCP throughput is given by the IP level throughput. The difference between the IP throughput and the throughput achieved by TCP seems to be a good measure of the efficiency of TCP. Figure 3 to figure 5 show the throughput over a number of 570 kbytes and 30 kbyte long file transfers for sample examples of a good position, a bad position and an unacceptable position.

A sample good position (figure 3)

The throughput computed per individual file transfer for the TCP remains quite close to the throughput on the IP level. In the case of short file transfer the variability of the throughput slightly increases.

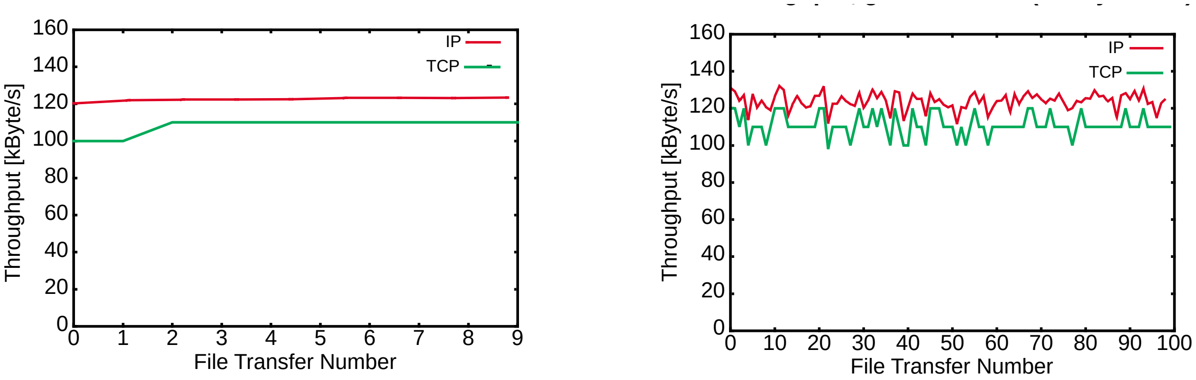

(a) (b)

Figure 3:Throughput (ARLAN) at a good position for (a) 570 kbyte file size (b) 30 kbyte file size

A sample bad position (figure 4)

The throughput computed on the IP level is - in the mean value - slightly lower than in the previous case, and - observed per individual file transfer - varies stronger, especially for 30 kbyte long files. The difference between the throughput on IP level and TCP level for transfers with 30 kbyte long files is sometimes small (less than 10 kbyte) and sometimes great (up to 40 Kbytes), the long term mean value at the transport level decreases significantly stronger for transfers with 570 kbytes long files.

A sample unacceptable position (figure 5)

The difference of the throughput on IP level and TCP level is always big. The short term variations of the IP level throughput seem to be mapped relatively close to the variations of the TCP throughput, in the way, that periods of very bad throughput and high throughput alternate with each other.

As we have mentioned the variability of the TCP throughput, it might be interesting to look at the plot of the standard deviation of the TCP mean throughput (figure 6) calculated for each location of the radio cell. The differences in the standard deviation are very high. One might say, that bad positions correspond to “hills”, unacceptable positions to “mountains” in fig 6.

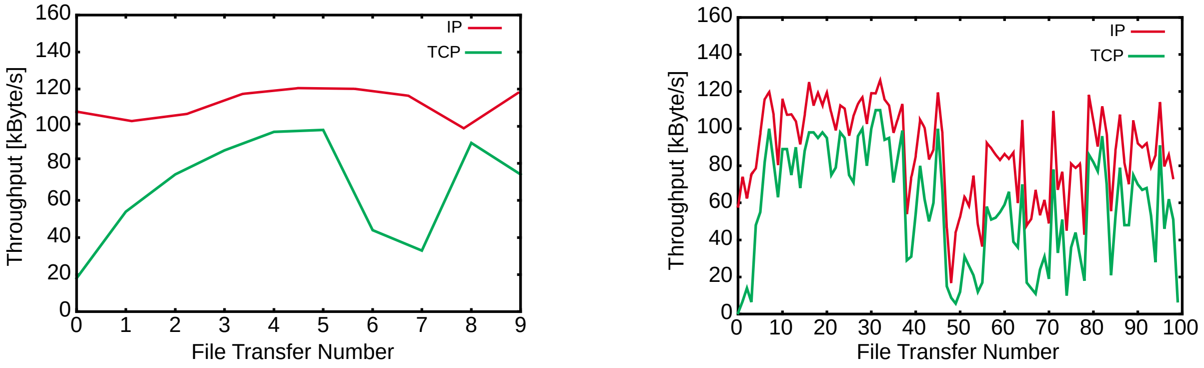

(a) (b)

Figure 4:Throughput (ARLAN) at a bad position for (a) 570 kbyte file size (b) 30 kbyte file size

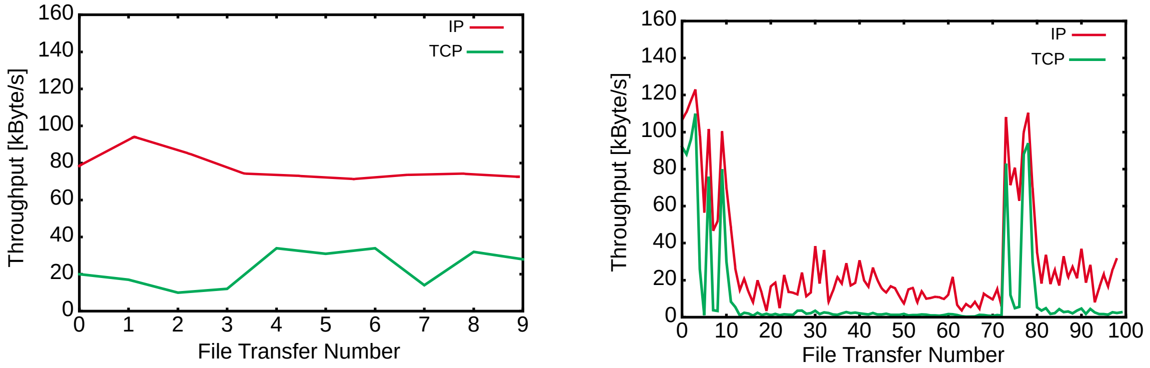

(a) (b)

Figure 5:Throughput (ARLAN) at an unacceptable position for (a) 570 kbyte file size (b) 30 kbyte f

The results of our experiments could be summarized as follows: The bad positions are all over the radio cell, also near the base-station. It is sometimes possible to find a bad position only at a distance of 2 meters from the base-station. The probability to put your working desk at a bad position is not negligible. In absence of this phenomenon one could simply suggest placing another base-station in one of the rooms 127 to 129 (which anyway would not solve the problem within the corridor, and cause overlapping of radio waves, at least as long as different radio channels are not used). Due to this phenomenon we have to state that although base-station placement is definitely an important question, even with reasonably increased number of base-stations the problem of throughput dependence upon location would still remain. Even in the "good" position the IP throughput is low compared to the channel bit rate of 2 Mbit/s. The potential for efficiency improvements might be expected within the physical layer and MAC Layer, but this is outside the scope of this paper. On the other hand there is very small potential for efficiency improvement in the transport level. As for the bad and unacceptable position, there is definitely a discrepancy between the IP- level throughput and the transport-level throughput.

The results obtained are similar using both ARLAN or WaveLAN. However, the distribution of good, bad and unacceptable positions is different for both technologies, due to their different implementations and their different protocol mechanisms. The size of a single radio cell of WaveLAN is greater compared to ARLAN, which is an indicator that the transmitter power and/or the receiver sensitivity of WaveLAN is higher. This leads to a slightly more stabilized throughput using WaveLAN, meaning that the "mountains" are less rugged compared to ARLAN. However, good, bad and unacceptable positions are arbitrary distributed using both technologies.

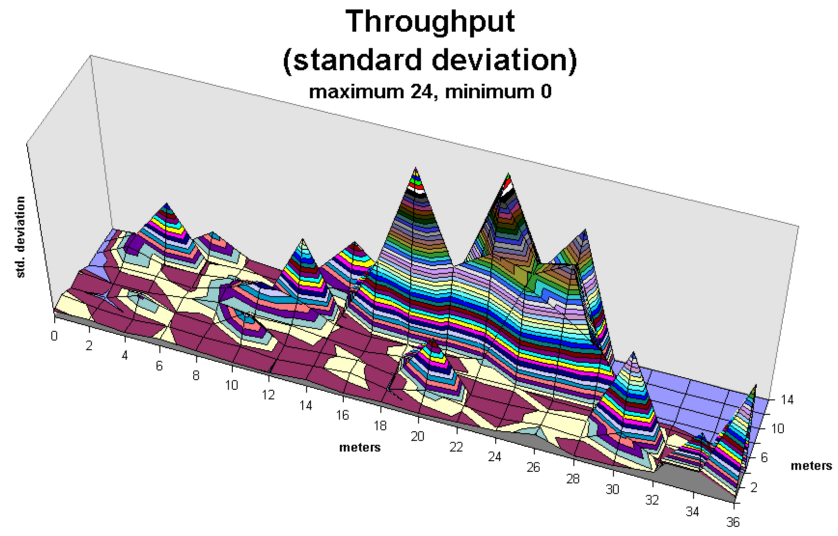

Figure 6:Standard deviation of the throughput for ARLAN

3.2. Performance at selected position

To get more insight into the operation of TCP, we investigated its reaction on the measured packet loss rate and the measured packet delay on IP level. To show TCP's operation we have measured TCP's congestion window, TCP's retransmission time-out value (rto) and the number of retransmissions performed by TCP. Due to the different characteristics of the used WLANs, the investigations are done separately for ARLAN and WaveLAN.

ARLAN Investigations

In case of ARLAN we measured a residual packet loss rate of about 10^{-6} due to the immediate acknowledgment mechanism. Therefore losses are not the reason for the throughput degradation at bad and unacceptable positions. In spite of this fact we observed retransmissions of TCP. Since with a packet loss rate of 10^{-6} losses are negligible. Therefore these retransmissions are due to wrong retransmission timer estimations. To corroborate the truth of this statement, we show (figure 7 to figure9) the measured TCP retransmission timer values (rto) and the IP delay of a packet for the first 15000 packets of a sample good, bad and unacceptable position with ARLAN.

.

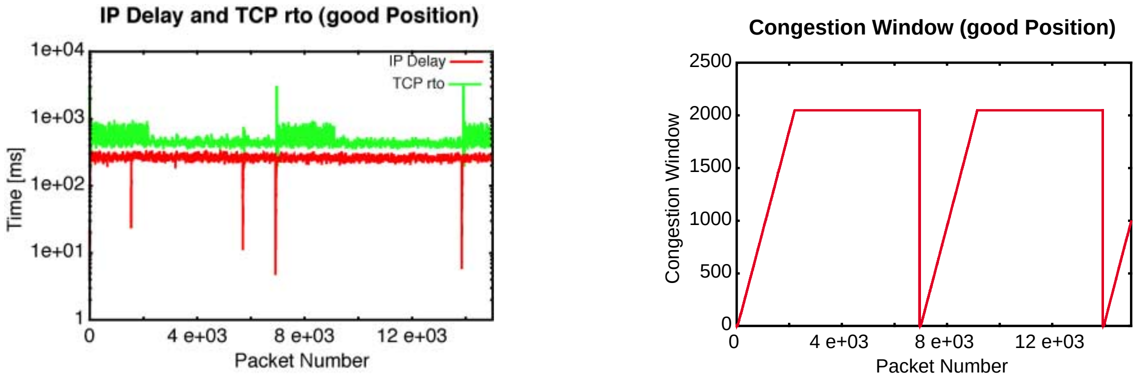

(a)

(b)

Figure 7:rto (a) and IP packet delay(b) at a good position for ARLAN

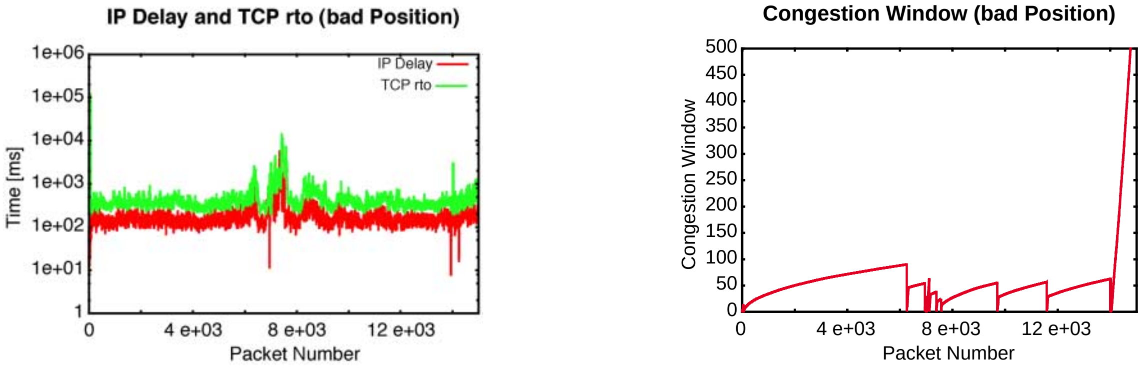

(a) (b)

Figure 8:rto (a) and IP packet delay(b) at a bad position for ARLAN

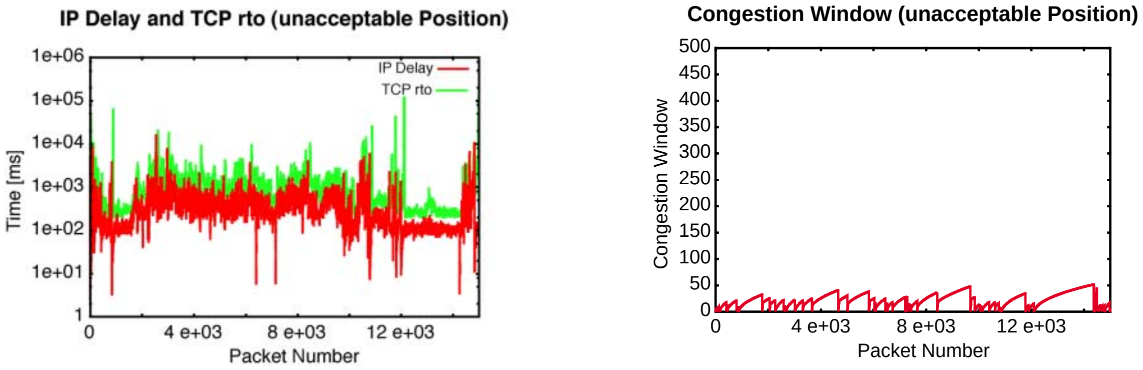

(a) (b)

Figure 9:r rto (a) and IP packet delay(b) at a unacceptable position for ARLAN

The measurements show, that in a good position the IP delay is nearly always smaller than the TCP rto (fig 7(a)). Comparing the IP delay and the TCP rto of a good position with a bad position (figure 8(a)), areas of overlapping curves are between the packets 6000 and 8000. In this area the IP delay is very often greater than the TCP rto, which causes unnecessary retransmissions. For the packets up to 6000 and the packets greater 8000, the IP delay is greater than the TCP rto like in the good position. The overlapping of the two curves is extreme within the unacceptable position (figure 9(a)). Over the range of 15000 packets the TCP rto is often smaller than the IP delay. The corresponding congestion windows are presented in figure 7(b) to figure 9(b) for each position The overlapping areas of figure 7(a) to figure 9(a) are corresponding with areas, where the congestion window (cwnd) decreases. At a good position (figure 7(b)) cwnd always increases. For a bad position (figure 8(b)), cwnd will be sometimes one which indicates a retransmission time-out. Looking at the unacceptable position (figure 9(b)), we note that very frequently cwnd becomes equal to one, being a good indicator for frequent retransmissions. That shows that TCP cannot adapt to the changing delay within a specific position. The resulting number of TCP retransmissions caused by the wrong rto estimation are shown in figure 10.

In case of the good position the TCP retransmissions are negligible, whereas for bad positions the retransmission rate is harmful and will increase dramatically in unacceptable positions, even under low losses.

Let us discuss in detail how the estimation of the time-out value depends on TCP's behavior: The sliding window approach of TCP in a single hop context means that multiple TCP segments will be buffered in the end-system waiting for transmission. For all waiting packets within a window the RTO value is equal. The delay of the nth packet to send, which resides in the buffer, depends on the delay of the TCP segments already waiting for transmission. Due to MAC retransmissions for each waiting packet the delay and its variation of the nth packet may be greater than the estimated RTO (including some overhead of 4*rttvar). In such a case TCP will produce unnecessary retransmissions and degrade throughput. Especially looking at Linux, where the granularity of RTO is 10ms, which is very small compared to the BSD implementation, a delay variation of the MAC is not so unusual.

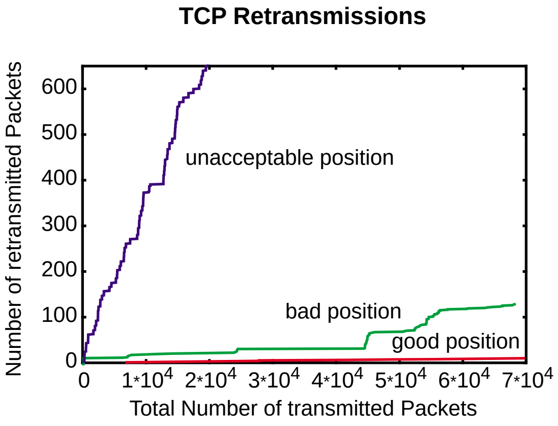

Figure 10:TCP retransmissions for ARLAN

WaveLAN Investigations

In case of WaveLAN the poor performance results were due to the error rate. At a good position the packet loss rate is around 10^{-6}, which is equal to the loss rate measured using ARLAN. At a bad position the packet loss rate grows up to 5% and at an unacceptable position the packet loss rate increases up to 40%. The reaction on losses is well known and described in the previous chapters: TCP interprets every loss as a sign of congestion and reduces the transmission rate.

4. Conclusions

The measurements confirmed that TCP is not able to support sufficient wireless Internet access. Even a reliable MAC is not able to overcome the poor TCP performance due to the great delay variability. Thus, similar results must be expected for 802.11 wireless LANs.

In addition the results show that it is necessary to consider the whole radio cell in order to investigate the efficiency of solutions. In particular it is desirable to minimize bad and unacceptable positions. Since we classified the achievable throughput into three types investigation of further approaches need not to redo the whole time consuming systematic measurement but only at specific positions.

5. Acknowledgments

We would like to thank T. Assimakopoulos for supervising the implementation of the bridge and Snoop as well as for many valuable discussions during the design and implementation. We also thank Enno Ewers who implemented the bridge and the Snoop extension within the scope of his minor thesis at our wireless lab. Further we thank Rolf Morich and Gerrit Schulte for their contribution to the development of the measurement environment.

6. References

[1] H Balakrishnan, V. Padmanabhan, S. Seshan, R. H. Katz: "A Comparison of Mechanisms for Improving TCP Performance over Wireless Links", IEEE/ACM Transactions on Networking, December 1997.

[2] H Balakrishnan, S. Seshan, E. Amir, R. H. Katz: “Improving TCP/IP Performance over Wireless Networks”, Proceedings of Mobicom, November ‘95

[3] ETSI Secretariat, “Hiperlan Functional Specification”, ETS 300 652, May 1997

[4] The editors of IEEE 802.11: "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specification", Draft Standard IEEE 802.11, P802.11/D5.0, July 96

[5] B. Rathke, T. Assimakopoulos, R. Morich, G. Schulte, A. Wolisz: “SNUFFLE: Integrated Measurement and Analysis Tool for Wireless Internet and its Use for in-House Environment”, 10th Intl. Conf. f. Computer Performance Evaluation, TOOLS’98, Sep. 98, Palma de Mallorca, Spain

[6] J. Weinmiller, M. Schläger, A. Festag, A. Wolisz: "Performance study of access control in wireless LANs - IEEE 802.11 DFWMAC and ETSI RES 10 HIPERLAN", Mobile Networks and Applications 2, 1997, pp 55-67

[7] M. Zitterbart, B. Stiller, A.N. Tantawy: “A Model for Flexible High-Performance Communication Subsystems”, IEEE Journal on Selected Areas in Communication, Vol. 11, No. 4, 1993

[8] M. Zorzi, R.R. Rao: “Energy constrained error control for wireless channels”, IEEE Personal Communications, vol. 4, pp. 27-33, Dec ‘97

[9] AT&T WaveLAN: PCM/CIA Adapter Card, http://www.wavelan.com

[10] Aironet PCM/CIA Adapter Card, http://www.aironet.com