

A Measurement Study on 802.11 Concurrently Used for Positioning and Communications

Thomas King, Thomas Haenselmann, Stephan Kopf, Wolfgang Effelsberg
{king,haenselmann,kopf,effelsberg}@informatik.uni-mannheim.de
Department of Computer Science
University of Mannheim
Germany

Abstract—Not only the communication capabilities of 802.11, but also the capability to determine the position of mobile devices make 802.11 highly appealing for many application areas. Typically, a mobile device that wants to know its position regularly performs active or passive scans to obtain signal strength measurements of neighboring access points. Active and passive scanning are survey techniques originally intended to be performed once in a while to learn about the presence and signal reception quality of access points within communication range. However, so far, no investigations are known to have been launched into how regular scanning affects concurrent data transmissions from an end-user point of view. In this paper, we explore how common data communication is affected while actively or passively scanning at the same time. We found that with an active scanning interval of equal or greater than 2 seconds the network conditions such as throughput and round trip delay are sufficient for interactive network applications. The use of passive scanning is prohibitive while simultaneously performing interactive data communication due to communication dropouts of more than 1.3 seconds during each scan.

I. INTRODUCTION

During recent years we have seen considerable improvements in downsizing computer hardware and in increasing the capacity of rechargeable batteries, as well as the advent of wireless networks for the mass markets. These technologies allowed manufacturers to build mobile devices that have a similar performance as desktop computers had several years ago. The benefit of mobile devices can be leveraged by so-called *location-based services*: Applications that act differently depending on the location of the user or, even better, proactively offer location-dependent information to the user. Location-based services are currently a hot topic in research, and are considered to be a promising market.

Nowadays, the *Global Positioning System* [1] is the predominant outdoor positioning system. Whereas GPS works well in many outdoor scenarios, it suffers from obstacles such as skyscrapers creating shielded street canyons or walls blocking the radio signal. To this end, a large number of research projects conceived novel positioning techniques (e.g. [2], [3], [4], and [5]). However, all these systems either require specialized hardware or show poor positioning accuracy.

Many recent research activities focus on *IEEE* 802.11-based positioning because almost everywhere, especially in occupied areas of developed countries, 802.11 network infrastructure is available for data communication [6]. Universities, offices and many private homes utilize 802.11 networks to get rid of

wires. As a reaction to the proliferation of 802.11, almost all modern mobile devices, ranging from smartphones to laptops, are shipped with built-in 802.11 network interfaces. 802.11 networks are not only used in indoor scenarios; even outdoors, many universities and coffee shops support nomadic users. Furthermore, 802.11-based positioning systems show sufficient positioning accuracy to be useful for a wide range of applications.

Another key argument for 802.11-based positioning systems is that 802.11 hardware can be used in dual mode: For data communication and for measuring the signal strength of neighboring access points. For instance, the *GUIDE* tourist guide [7] and *PlaceLab* [8] are two representatives for outdoor positioning systems that utilize 802.11 for data transmissions as well as position determination. Many indoor positioning systems such as *RADAR* [9], *HORUS* [10] and *COMPASS* [11] require a high rate of signal strength readings especially if they are running in tracking mode.

In 802.11 the typical way of measuring the signal strength of access points within communication range is to perform active or passive scans. Active and passive scanning are survey techniques originally intended to be performed once in a while to learn about the presence and signal reception quality of access points within communication range. So far, it was unknown what happens to the data communication capabilities of mobile devices if signal strength measurements in form of active or passive scans are performed concurrently. In our preliminary work [12] we briefly discussed that a high rate of signal strength measurements affect the data communication capabilities of 802.11 network cards. This work goes a step further and provides an in-depth investigation how throughput and round trip delay suffer from different scan techniques and intervals. The results of this work can be used as guidance to select an appropriate scanning technique and scanning interval for 802.11-based positioning systems while at the same time enable mobile devices to perform meaningful interactive data communication.

The rest of the paper is organized as follows. The subsequent section (Section II) presents the relevant related work. Section III introduces active and passive scanning. In Section IV, we discuss how active and passive scanning affect the communication capabilities of mobile devices. Finally, a conclusion is provided in Section V.

II. RELATED WORK

Several previously published studies investigate throughput and delay on 802.11. Xylomenos and Polyzos [13] explore the throughput of UDP and TCP achievable with several early 802.11 network cards. Their research focuses on throughput limitations caused by software implementation issues. The researchers recommend changes in the implementations of network protocols as well as a modification of drivers.

Duchamp and Reynolds measure throughput while varying the distance between a mobile device and an access point [14]. In [15], Bing measures delay and throughput for two early 802.11 network interfaces in a lab environment. A performance degradation is observed by Heusse et al. [16] if some mobile devices use a lower bit rate than other devices. The authors analyze the problem theoretically as well as empirically and derive a simple expression for the throughput.

Compared to our work, all these approaches do not consider scanning at all and use rather out-dated 802.11 or 802.11b hardware.

Many research projects exist in the area of 802.11-based positioning systems. Bahl et al. are the first to investigate 802.11-based positioning systems [9]. Although, they talk about driver difficulties in early versions of 802.11 hardware, they never covered the problem of using 802.11 hardware simultaneously for communications and regular scanning in their papers.

Researchers at Rice University undertook two research projects to investigate 802.11-based positioning systems [17]. Both projects focus on algorithms to compute the location of a mobile device while being carried around by a user. While they figured out that some 802.11 network card drivers report the same access point more than once during a scan, they have not worked on the impact of regular scanning on data communication capabilities of 802.11 hardware.

One of the largest research projects in the area of positioning systems in the last few years is PlaceLab [8]. PlaceLab focuses not only on 802.11-based systems, instead it utilizes a wide range of already existing Beacons to compute a position estimate (e.g., Bluetooth [8], GSM [18]). Because PlaceLab covers so many sources of Beacons, they never dug into specific details of 802.11 related problems.

III. ACTIVE VS. PASSIVE SCANNING

In this section, we discuss two techniques to discover neighboring access points and to measure their signal strength, namely *active scanning* and *passive scanning*, as described by the IEEE 802.11 standard [19].

A. Functional Principles

IEEE 802.11 subdivides the radio spectrum into a set of channels. The number of available channels depends on where 802.11 is used and which sub-specification of the 802.11 physical layer is selected. For instance, in the United States, only 11 channels are allowed for 802.11b and 802.11g, whereas 13 channels can be used in Europe. In contrast, the commercially less successful 802.11a defines 12 channels, however, in some countries the radio spectrum of 802.11a is still assigned to other purposes, today.

A wireless network interface listens to one channel at a given time. So, if a mobile station wants to get to know all the access points in communication range, it has to tune its wireless network interface to each channel, one after another, and perform a scan.

IEEE 802.11 defines two scanning techniques: Active scanning and passive scanning. The former approach requires a bi-directional communication initiated by the mobile station. For the latter approach, the mobile station passively listens for management frames sent out by access points. The details of these two techniques are discussed in the following two sections.

For the remainder of this paper, we focus on the infrastructure mode of 802.11 because this is the typical scenario in the field of positioning systems. We mainly focus on 802.11g because it is the latest sub-standard of the 802.11 standard family and it is one of the most frequently used ones. However, our results are also applicable to 802.11b. Furthermore, our scenario is located in Europe and 802.11 operates on 13 channels.

B. Active Scanning

A mobile device follows the subsequent procedure for each channel to perform an active scan: It tunes the wireless interface to the particular channel. The mobile device waits for either incoming frames generated by other devices or for a timer to expire. The timer makes sure that the mobile device is only waiting a certain period of time for incoming frames. After that, it uses the 802.11 medium access procedure to gain access to the channel and sends a so-called *Probe Request* frame. Access points that receive a Probe Request frame are supposed to reply with a so-called *Probe Response* frame. The mobile device waits for a certain time, and if no frame is received it proceeds to the next channel. If any Probe Response frame is received, the mobile device processes it for further use. A Probe Response frame contains information such as supported data rates, the name of the network, and the access point's MAC address. By examining received Probe Responses, a mobile device is able to recognize neighboring access points and their signal strength.

The IEEE 802.11 standard does not define default values for the aforementioned timers, however, [20] and [21] empirically studied the values used by wireless network card manufactures. In total, the exact time required to perform an active scan can vary significantly based on the number of available access points and hardware capabilities. In our measurements, we found that an active scan over all channels takes less than 260 milliseconds to be completed.

C. Passive Scanning

Passive scanning has been introduced to reduce the workload of mobile devices and hence save battery power. While scanning passively, a device listens to each channel and waits for a given period of time. If an access point is assigned to a particular channel, the mobile station should receive a so-called *Beacon* frame. Every access point broadcasts Beacon frames on a regular basis to maintain the network. Beacons usually contain the same information as Probe Response frames, such as supported data rates, the name of the network, and the MAC address of the access point. By examining the received Beacon frames, a

mobile device is able to recognize neighboring access points and their signal strength.

Once in a while, an access point sends out a Beacon frame that carries additional management information: A *Delivery Traffic Information Map* (DTIM). The idea is that battery powered mobile devices may sleep in low power mode while the access point buffers frames for these stations. The DTIM indicates which stations have buffered traffic waiting to be picked up and hence this map should be received by every mobile device associated to a given access point. This means that a mobile device should never miss such a Beacon regardless if it is sleeping or scanning. In case a scan is performed while a DTIM is scheduled by the access point the mobile device is associated with, it is forced to cancel the scan process and wait for the map. For this, Beacons contain the beacon interval as well as how many regular Beacons will be transmitted before the next DTIM. This information allows a mobile device to determine the point in time when the next DTIM will be transmitted by the access point it is associated with. The rate of DTIMs is often configurable, and commonly every 10^{th} Beacon is used for this purpose.

Access points usually broadcast a Beacon packet every 100 milliseconds which means that a mobile station should stay on a particular channel at least for the same period of time to make sure not to lose a Beacon from an unknown access point. In total, a passive scan requires at least 1.3 seconds to be completed. This is nearly five times the length required for an active scan.

IV. EFFECTS ON COMMUNICATION PERFORMANCE

In this subsection, we investigate how active and passive scanning affect regular data transmission.

A. Experimental Environment

To achieve interpretable results we simplified our environment: One mobile device communicates with one access point. In a more complex scenario with additional mobile devices, throughput and round trip time may even be worse and more volatile.

We used a *Fujitsu Siemens* Lifebook T4010 laptop running *Linux* kernel 2.6.16 and *Wireless Tools* 28pre13. We implemented passive scanning support into the ipw2200 version 1.1.3 network interface driver [22], so that we were able to use the built-in *Intel PRO/Wireless 802.11b/g* network card of the laptop.

A *Linksys / Cisco* WRT54GS access point assigned to channel 8 has been used to gain access to the local network of the *University of Mannheim*. The WRT54GS access point was running the *Alchemy* firmware version 1.0 and was configured for 802.11b/g mode with a beacon interval of 100 milliseconds and a DTIM every 10^{th} Beacon. The distance between the laptop and the access point was approximately 3 meters, and during the measurements a 54 MBit/s link between the laptop and the access point was established.

We conducted data transmission measurements with *iperf* version 2.0.2 [23] to gauge throughput and round trip time. For this, we used an *iperf* server within the local network

and an *iperf* client running on the laptop. The *iperf* server was connected to the access point via a 100 MBit/s switched Ethernet, so that the wireless link was the only bottleneck. *Iperf* was configured to measure the throughput every 0.5 seconds and to transmit data for 60 seconds. For all the graphs presented in this paper, we carried out the experiments at least three times and selected the result showing the highest throughput.

B. Experimental Results

Throughput and round trip delay are the main objectives and are first measured without any scanning at all to get a reference. Based on this reference, throughput and round trip time are quantified for various scan intervals and different scan schemes. The relation between the maximum throughput and the throughput achievable for a particular scan interval gives a well-balanced estimate on how common data communication is affected by scanning. Additionally, the round trip time measurements indicate how interactive data communication is strained by concurrent scans. Especially, network applications that are interactively used by users suffer from high round trip delays. For instance, video conferencing applications or Internet telephony applications as well as simple remote shells such as SSH or Telnet are only usable if the round trip delay is under a certain level.

From [24] we know that a round trip delay below 290 milliseconds between user input and reaction of the application is usually tolerated by humans. Telephony applications such as Skype or SIP require similar delay boundaries between mouth and ear to deliver acceptable quality [25]. Based on these values, we define a network as *stable* if the round trip delay is below 300 milliseconds in 90 percent of all measurements.

1) *UDP*: First of all, we measure the maximum throughput that is achievable with UDP over an 802.11g link. For this, we stepwise increase the bandwidth acquired by our experiments from 15 MBit/s to 30 MBit/s. Our results show that the throughput reaches the peak at a sending data rate of 27 MBit/s. Figure 1 shows the achieved throughput and round trip time. As we see from this figure, the throughput varies between 22.25 and 26.75 Mbit/s while the round trip time is between 3 and 323 milliseconds. On average, the throughput is 25.155 MBit/s and the round trip delay is 82.30 milliseconds. The variations in throughput and delay are due to retransmissions on the MAC layer and changing radio channel characteristics.




Figure 1. Throughput and round trip time of UDP data transmission over an 802.11g link.

In our next experiment, we perform active scans at a high rate to meet the requirements of tracking systems. Our wireless network card requires at least 260 milliseconds to perform an active scan, so we set the scan interval to 0.3 seconds.

This means that only 40 milliseconds can be used for data transmissions between two consecutive scans. As we see from Figure 2(a), the throughput does not exceed 7 MBit/s and is very often lower than 1 MBit/s, resulting in an average throughput of 797 KBit/s with a standard deviation of 1.571 MBit/s. In other words, only 3 percent of the throughput is available in comparison to performing no scanning at all. The throughput peaks around seconds 6, 13, 22, 45, and 49 are due to canceled scans. The reason for this is that a Beacon carrying a DTIM is scheduled by the access point as mentioned in Section III-C.

Furthermore, if we look at the round trip time we see that it varies between 23.5 and 6656 milliseconds resulting in an average delay of 2400 milliseconds with a standard deviation of 1331 milliseconds. Because average and standard deviation are sometimes misleading, we present the cumulative distribution function of the round trip delay in Figure 2(b). From this distribution we see that in 80 percent of all measurements the delay is larger than 900 milliseconds. Even worse, in more than 10 percent of all cases the round trip time is larger than 4 seconds.

With such a high round trip time in conjunction with low throughput, meaningful communication is often not feasible. Especially, if we consider typical interactive network applications such a network condition cannot be considered stable.

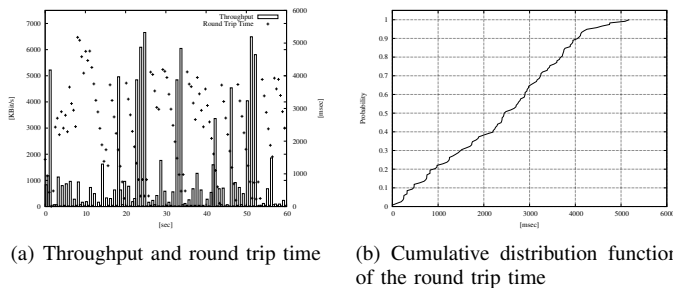



Figure 2. Throughput and round trip time of UDP data transmission while actively scanning every 0.3 seconds.

Our next experiment shows the trade-off between active scan intervals versus throughput and round trip delay. For this, we carried out measurements with active scan intervals ranging from 0.3 seconds to 4 seconds. Figure 3(a) depicts the results. As we see from the figure, an active scan interval of 0.5 seconds still produces a round trip delay of nearly 500 milliseconds and limits the throughput to 3.5 MBit/s. However, an active scan interval equal to or greater than 2 seconds produces acceptable results. On average, a throughput of more than 20 MBit/s (this corresponds to nearly 80 percent of the throughput that is available if no scanning is performed) and a round trip delay of less than 150 milliseconds is achievable. Figure 3(b) depicts the cumulative distribution function for the round trip delay to provide deeper insights in the delay distribution. As we see from this figure, a scan interval of 2 seconds generates in nearly 90 percent of all cases a delay of less than 300 milliseconds. With a scan interval of 4 seconds, the 90th percentile is less than 180 milliseconds.

Furthermore, if we compare the two curves generated by the two scan intervals we see that the first parts of both curves are

more or less identical. The reason for this is that the packets used for these delay measurements get through without any interference caused by scanning. We see from the figure that more than 70 percent of all packets are transferred without any major impact caused by scanning. On the other hand we see less than 2 percent of all measurements are twice affected by scanning. Measurements that show a round trip time of more than 500 milliseconds are delayed on the way from the mobile device to the iperf server, and are additionally buffered on the access point on the way back from the iperf server to the client.

Based on these observations we summarize that an active scan interval of equal or greater 2 seconds renders network conditions well enough to allow meaningful communication. Furthermore, dependent on the positioning system used and network applications running, an active scan interval can be selected to fit the needs of both applications.



(a) Average and standard deviation of throughput and round trip time (b) Cumulative distribution function of the round trip delay for scan intervals of 2 and 4 seconds

Figure 3. UDP throughput and round trip time while concurrently performing active scans with different scan intervals.

In addition, we also investigate how passive scanning affects concurrent data transmission. While an active scan can be completed in less than 300 milliseconds, a passive scan requires at least 1300 milliseconds to terminate. For instance, the wireless network card we use stays 120 milliseconds on each channel while performing a passive scan, resulting in an overall passive scan time of 1560 milliseconds. Being busy with scanning such a long time may dramatically limit the achievable throughput. To investigate this phenomenon, we carried out experiments with a passive scan interval between 2 and 9 seconds. As shown by Figure 4(a), with a passive scan interval of 2 seconds only an average throughput of slightly more than 5 MBit/s is achievable while the round trip delay is approximately 700 milliseconds. To get a deeper understanding of the delay distribution we plotted the cumulative distribution function in Figure 4(b). From this figure we see that in 30 percent of all cases the round trip delay is larger than 1 second if a scan interval of 2 seconds is applied. Such a high round trip delay makes interactive data communication nearly impossible.

As we see from Figure 4(a), it takes at least a scan interval of 9 seconds to obtain an average throughput and round trip delay that is sufficient for common data transmissions. While with a scan interval of 9 seconds the average throughput reaches 19 MBit/s, a problem still arises from the round trip delay. Figure 4(b) depicts the cumulative distribution function of the round trip delay for the 9 seconds scan interval. As we see from the figure, more than 80 percent of all measurements

show a delay of less than 200 milliseconds. However, more than 5 percent of all packets require more 1 second to travel round trip from the client to the server and back. The tail is far longer than the tail obtained from active scans. The reason for this is that passive scanning stops communications for more than 1300 milliseconds during a scan. Based on these results, we summarize that passive scanning is inappropriate if it is used in combination with interactive data communication via UDP. Furthermore, compared to the active scanning approach, we see that passive scanning disturbs throughput and round trip delay even more.

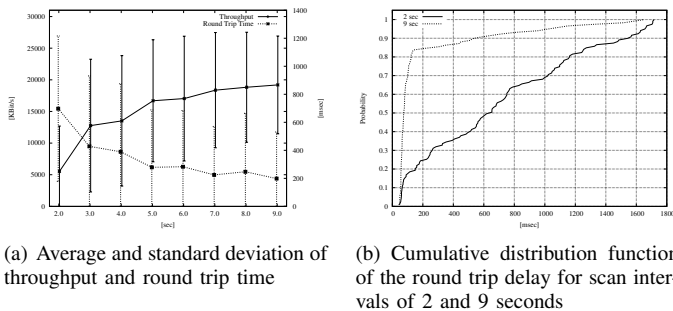


Figure 4. UDP throughput and round trip time while passively scanning with different scan intervals.

2) *TCP*: In this paragraph, we analyze how scanning affects TCP data transmissions. We distinguish between UDP and TCP traffic because UDP does not have any congestion control mechanisms and hence in our simple scenario it shows the maximum achievable throughput. On the other side, TCP is the most frequently used transport protocol, and its congestion control algorithms might be confused about communication dropouts caused by frequent scanning operations. Especially, the large and abrupt delay changes we have observed in our UDP study might confuse the congestion control algorithm. The congestion control mechanism of TCP utilizes round trip time measurements to determine if the network gets congested or if it is already clogged. Additionally, the large buffers utilized by TCP to control congestion and flow might add another source of delay.

To get a reference value of how much data can be transferred over a 802.11g link, we invoked iperf in TCP mode and sampled the throughput and round trip delay for 60 seconds. Our measurements show, on average, a throughput of 17.1 MBit/s and a round trip time of 39.3 milliseconds; the standard deviation of the measurements is 768 KBit/s for throughput and 11.82 milliseconds for delay. We see a large difference between TCP and UDP throughput; this difference has already been studied by George Xylomenos et al. in [13].

To make the results of the TCP measurements comparable to the ones obtained for UDP, we used the same scan intervals. Figure 5(a) shows average throughput and average round trip time as well as standard deviations for both measures in the active scanning scenario. An active scan interval of 0.3 seconds results in 547 KBit/s and a round trip time of 1700 milliseconds. 547 KBit/s correspond to nearly 3 percent of the maximum achievable TCP throughput which is comparable to the UDP

results. The average round trip delay of TCP is only 70 percent of the UDP delay. The reason for this is that TCP does not fill the channel in a way UDP fills it, because after a scan is completed the congestion control mechanism of TCP needs some time to take advantage of the available bandwidth. This results in less filled buffers in the network stack and hence in a smaller round trip delay.

An active scan interval of 1 second and TCP as transport protocol produce 10.4 MBit/s throughput. In other words, more than 60 percent of the throughput can be obtained in comparison to no scanning. The round trip delay also shows relatively stable values around 90 milliseconds with a small standard deviation of nearly 60 milliseconds. Figure 5(b) depicts the cumulative distribution function. This figure shows due to the less filled buffers and a large share of spare bandwidth that with a scan interval of 1 second the delay is always under 265 milliseconds. These are interesting results, especially in comparison to the UDP results, where stable network conditions can be observed only with an active scan interval larger or equal to 2 seconds. Even worse, with a scan interval of 4 seconds and TCP as transport protocol the tail of the delay distribution grows. The reason for this is, as already mentioned, the use of larger buffers by the congestion control mechanism of TCP. These increased buffers are required to fill the anticipated bandwidth.

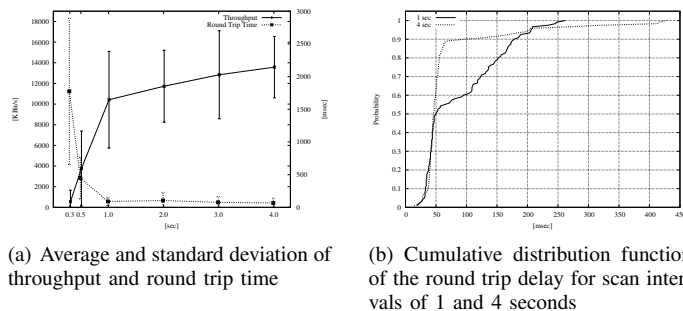


Figure 5. TCP throughput and round trip time while performing active scans with different scan intervals.

In the following experiment, we compare TCP and UDP data transmission while performing passive scans. With a passive scan interval of 2 seconds TCP nearly achieves 2.2 MBit/s; that is less than 35 percent of what is achievable with UDP (see Figure 6(a)). If we further compare the throughput of TCP and UDP, we see that the ratio nearly converges against equality if the scan interval is increased. For instance, for a scan interval of 5 seconds the ratio is 1:2 and a scan interval of 9 seconds shows a ratio of 1:1.27. The reason for this is, as aforementioned, that the TCP congestion control mechanism is confused by the large and abrupt changes in the delay. If the time between two subsequent scans is increased, the congestion control mechanism is able to adapt to the situation and fill up the spare bandwidth.

The round trip delay shows a similar pattern for TCP as for UDP. As Figure 6(b) shows, with a scan interval of 2 seconds, more than 40 percent of all measurements need more than 1000 milliseconds to be completed. Even worse, the 90th percentile shows a delay of more than 2250 milliseconds. This leads to the fact that a scan interval of 2 seconds is prohibitive if

interactive data communication is simultaneously performed. A scan interval of nine seconds still leads to an average round trip time of 128 milliseconds and a huge standard deviation of more than 246 milliseconds. As the cumulative distribution function of Figure 6(b) shows, result more than 8 percent of all measurements in a delay of more than 500 milliseconds. Even more than 5 percent result in a delay of more than 750 milliseconds. These results show that passive scanning is also inappropriate if concurrently interactive data communication via TCP are performed.

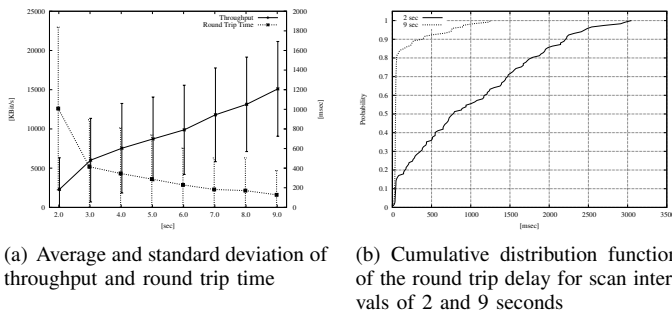


Figure 6. TCP throughput and round trip time while concurrently performing passive scans with different scan intervals.

To sum up, from this section we see that passive scanning is inadequate if concurrent interactive data transmissions are performed. This fact is regardless of the transport protocol used and the scanning interval applied. The reason for this is that passive scanning generates a communication dropout for more than 1300 milliseconds each time a scan is performed.

Active scanning in contrast can be used simultaneously while running interactive network applications if a scanning interval of equal or greater 2 seconds is applied. In this case, the network conditions can be considered stable enough to allow meaningful interactive data communication.

V. CONCLUSIONS

The contribution of this paper is an empirical analysis of how active and passive scanning affect concurrent data transmission. We first introduced the concept of scanning and then briefly summarized how active and passive scanning works. Subsequently, we provided an empirical analysis of how active and passive scanning interferes with common interactive data communication. For this, we investigated the two predominant transport protocols TCP and UDP. We found out that with an active scanning interval of 2 seconds, the resulting network conditions can be considered stable enough so that meaningful interactive network applications can be concurrently run. Furthermore, we found that the use of passive scanning is prohibitive while simultaneously performing interactive data communication. The reason for this is a communication dropout of more than 1.3 seconds each time a scan is executed.

ACKNOWLEDGMENTS

The authors acknowledge the financial support granted by the *Deutschen Forschungsgemeinschaft* (DFG).

REFERENCES

- [1] Kaplan, E., Hegarty, C., eds.: Understanding GPS: Principles and Applications. Second edn. Artech House Incorporated (December 2005)
- [2] Priyantha, N.B., Chakraborty, A., Balakrishnan, H.: The Cricket Location-Support System. In: Proc. of the 6th ACM MobiCom. (2000) 32–43
- [3] Ni, L.M., Liu, Y., Lau, Y.C., Patil, A.P.: LANDMARC: indoor location sensing using active RFID. *Wireless Networks* **10** (November 2004) 701–710
- [4] Want, R., Hopper, A., Falcao, V., Gibbons, J.: The Active Badge Location System. *ACM Transactions on Information Systems* **10**(1) (January 1992) 91–102
- [5] Brumitt, B., Meyers, B., Krumm, J., Kern, A., Shafer, S.: EasyLiving: Technologies for Intelligent Environments. In: Proc. of the 2nd International Symposium on Handheld and Ubiquitous Computing. (2000) 12–27
- [6] arkasha and bobzilla: Wireless Geographic Logging Engine - Plotting WiFi on Maps. Website: <http://www.wigle.net> (2001-2007)
- [7] Cheverst, K., Davies, N., Mitchell, K., Friday, A.: Experiences of Developing and Deploying a Context Aware Tourist Guide: The GUIDE Project. In: Proc. of the 6th ACM MobiCom. (2000) 20–31
- [8] LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I., Scott, J., Sohn, T., Howard, J., Hughes, J., Potter, F., Tabert, J., Powledge, P., Borriello, G., Schilit, B.: Place Lab: Device Positioning Using Radio Beacons in the Wild. In: Proc. of the 3rd IEEE PerCom. (2005) 116–133
- [9] Bahl, P., Padmanabhan, V.N.: RADAR: An In-Building RF-Based User Location and Tracking System. In: Proc. of the 19th IEEE InfoCom. (2000) 775–784
- [10] Youssef, M.: Horus: A WLAN-Based Indoor Location Determination System. PhD thesis, University of Maryland at College Park (2004)
- [11] King, T., Kopf, S., Haenselmann, T., Lubberger, C., Effelsberg, W.: COMPASS: A Probabilistic Indoor Positioning System Based on 802.11 and Digital Compasses. In: Proc. of the 1st ACM WinTECH. (2006) 34–40
- [12] King, T., Haenselmann, T., Kopf, S., Effelsberg, W.: Overhearing the Wireless Interface for 802.11-based Positioning Systems. In: Proc. of the 5th IEEE PerCom. (2007) 145–150
- [13] Xylomenos, G., Polyzos, G.C.: TCP and UDP Performance over a Wireless LAN. In: Proc. of the 18th IEEE InfoCom. (1999) 439–446
- [14] Duchamp, D., Reynolds, N.F.: Measured Performance of a Wireless LAN. In: Proceedings of the 17th Conference on Local Computer Networks (LCN), Minneapolis, MN, USA, IEEE (September 1992) 494–499
- [15] Bing, B.: Measured Performance of the IEEE 802.11 Wireless LAN. In: Proc. of the 26th IEEE LCN. (1999) 34–42
- [16] Heusse, M., Rousseau, F., Berger-Sabbatel, G., Duda, A.: Performance Anomaly of 802.11b. In: Proc. of the 22nd IEEE InfoCom. (2003) 836–843
- [17] Haeberlen, A., Flannery, E., Ladd, A.M., Rudys, A., Wallach, D.S., Kavradi, L.E.: Practical Robust Localization over Large-Scale 802.11 Wireless Networks. In: Proc. of the 10th ACM MobiCom. (2004) 70–84
- [18] Varshavsky, A., LaMarca, A., Hightower, J., de Lara, E.: The SkyLoc Floor Localization System. In: Proc. of the 5th IEEE PerCom. (2007) 125–134
- [19] Institute for Electrical and Electronics Engineers, Inc.: ANSI/IEEE Standard 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Website: <http://standards.ieee.org/getieee802/> (1999)
- [20] Arbaugh, W., Shin, M., Mishra, A.: An Empirical Analysis of the IEEE 802.11 MAC Layer Handoff Process. *ACM SIGCOMM Computer Communications Review* **33**(2) (April 2003) 93–102
- [21] Velayos, H., Karlsson, G.: Techniques to Reduce the IEEE 802.11b Handoff Time. In: Proc. of the IEEE ICC. (2004) 3844–3848
- [22] King, T.: Passive Scanning not supported? ipw2100 developer mailinglist: <http://sf.net/mailarchive/forum.php?thread%5Fid=27023011%5F&forum%5Fid=38%938> (July 2006)
- [23] Gates, M., Warshavsky, A., Pietsch, J., Cerveny, B., Lambert, M., Finkelson, D., Zekauskas, M., Zekauskas, M., Oliver, K.: Iperf. Website: <http://dast.nlanr.net/Projects/Iperf> (August 2006)
- [24] Mauve, M., Vogel, J., Hilt, V., Effelsberg, W.: Local-lag and Timewarp: Providing Consistency for Replicated Continuous Applications. *IEEE Transactions on Multimedia* (February 2004) 47–57
- [25] Janssen, J., Vleeschauwer, D.D., Petit, G.H.: Delay and Distortion Bounds for Packetized Voice Calls of Traditional PSTN Quality. In: Proceedings of the First IP-Telephony Workshop (IPTel). (2000) 105–110