2016719210

A Comparison of Mechanisms for Improving TCP Performance over Wireless Links

Olgun, T.

Abstract—The packets that transfers via TCP protocols are in danger of loss and those losses are often the result of congestion. However the packet transfers over wireless networks often have the problem of being lost due to non-congestion related sources such as handoffs and bit-errors. The TCP protocols are mostly based on integrating a congestion prevention system when it comes to situations like this which results in less performant network. In this paper several different TCP approaches are used to determine the performance comparison in such wireless networks. The results show that the end-to-end connection without being splited at the base station brings good performance and so the selective acknowledgments and explicit loss notifications. Also, reliable link-layer protocol has a good performance over TCP.

***Index Terms*—E2E, End-to-end schemes, Link-layer schemes, TCP, TCP-Reno**

I. Introduction

In this study, various TCP mechanisms and some tests and results of those mechanisms are given according to their performance over wireless networks. The study is based on some questions that deemed to be required to be answered. The main idea of these questions are,

- Find the best performance for each combination of TPC mechanisms
- Understand the importance of link-layer schemes to be aware of so that it achieves high throughput
- Find the usefulness of selective acknowledgments in presence of burst losses
- Figure out the importance of split connection for sender to be guarded from wireless losses

II. Main Topics

As most of the protocols including TCP are designed for wired networks, these protocols assume congestion as the fundamental reason for packet losses or unconventional delays. To prevent a packet loss the congestion prevention system of TCP works as following: It takes a round trip delay time and calculates standard deviation. If no acknowledgment arrives or the timeout which has been calculated is passed then the TCP server realizes the problem and halve its congestion window due to its congestion control mechanisms [1, 2]. The problem here is that if the reason behind a packet loss or unusual delay which causes TCP server to activate avoidance mechanisms is other than congestion, it results in unnecessary slow down on the throughput .

There are two possible ways to solve this issue. The first is a mechanism that hides the effects from TCP sender so it does not require an additional tweaking for existing implementations. The second is to make the sender aware of there is a wireless hop so not all packet losses are caused by congestion.

A. End to end schemes

Though there are numerous TCP variants available and being used on the Internet, the most used variant of TCP is TCP-Reno. Adding selective acknowledgements to the standard allows the sender to handle losses more efficiently but the problem is that the sender still thinks the source of the losses is congestion and shrinks its congestion window size. Via this behavior the percentage of the end-to-end performance can be identified. For the E2E+SACK protocol, SMART (Simple Method to Aid ReTransmissions) approach has been used. The sender resends a packet if there is a SMART acknowledgment. E2E+ELN adds Explicit Loss Notification (ELN) to the protocol. The acknowledgments after packet loss that are non congestion related are modified to inform the sender about the packets so the sender may resend the packets without having congestion mechanisms open.

B. Link-Layer Schemes

Link layer protocols have no standard TCP variant as in transport layer protocols. The link layer algorithm that has been used in the research is called LL which with the help of cumulative acknowledgments, determine the packet losses. It uses existing TCP acknowledgments and the round-trip estimate is smoothed. This allows it to resend packets more than once before TCP-Reno would time out. For LL-SACK protocol, selective retransmission is used to improve the

This paper, a quick summary of A Comparison of Mechanisms for Improving TCP Performance over Wireless Links, 1996 was sent for review to Alagoz, F (fatih.alagoz@boun.edu.tr) and Gur, G (gurkan.gur@gmail.com) as a draft on 4th of Dec, 2017 and second version on 29th of Dec, 2017.

Olgun, T. studies Software Engineering at Bogazici University. He also works as a back-end developer for Ambeent Wireless A.S. at Istanbul Technical University Ari Teknokent building. (e-mail: tugcan.olgun@boun.edu.tr)

performance. LL-SACK works as SMART though it uses TCP acknowledgments instead of generating one.

C. Split-Connection Schemes

The implementation of split-connection scheme depends on avoiding the copying of data in the intermediate host by passing pointers. The SPLIT scheme that has been used here is like I-TCP, uses an intermediate host to divide a TCP connection into two separate TCP connections [1].

III. Experimental results

The test equipment consist of IBM ThinkPad laptops, Pentium PCs running BSD/OS2.0. The devices are connected with 10 Mbps ethernet and 915 MHz AT&T WaveLANs and shared wireless with bandwidth of 2 Mbps. TCP data transfers peak at 1.5 Mbps.

The TCP data packet sizes are 1400 bytes and the average error rate is one every 65 KBytes. Which means the bit-error rate is about 1.9e-6. WAN experiments carried away with 16 Internet hops to investigate the impact of large delay-bandwidth products.

In each run 8MBytes transfer from the source to receiver is conducted and this is repeated several times for consistency. Throughput at the receiver, wireless and wired goodputs, with the help of tcpdump the and congestion related variables are measured during the research.

With bit error rate of 1.9e-6, the packet error rate is 2.3%, wireless goodput is 97.7% and wired goodput is 100% if no congestion exists.

In summary simple link-layer retransmission scheme has the opposite effect on TCP performance while enhanced link-layer scheme has a way better performance on TCP. Furthermore, on the contrary of split-connection, TCP acknowledgments' end-to-end semantics are conserved in link-layer protocol. To get a good performance, end-to-end connection need not be split at the base station.

IV. Conclusion

In this paper, TCP protocols over wireless communication over several hops are investigated and their performances are measured based on the given criteria. They have been categorized as end-to-end, split-connection based and link-layer. End-to-end throughput over wired and wireless connection is used to conduct the research and do the comparisons accordingly.

Conclusions are:

1) On the contrary of a regular TCP connection a reliable link-layer protocol which is integrated with knowledge of TCP to shield gives 10-30% better throughput. Of the protocols and schemes that have been tested, the best performant is TCP-aware link-layer protocol.

2) Using SMART-based selective acknowledgments mechanisms over wireless communication gives better performance than the split-connection approach which usually ends up being stalled due to timeouts. Nevertheless the mechanism still has slightly worse performance than link-layer protocol which indicates it is not strictly necessary for achieving good performance.

3) Using SMART-based selective acknowledgments over LAN connection has very good performance when it comes to compensate the packet loss rate over hops.

4) Even though compared to local techniques in dealing with wireless losses, end-to-end schemes look promising to get a good performance.

References

[1] A. Bakre and B. R. Badrinath. (1995, May) I-TCP: Indirect TCP for Mobile Hosts. In Proc. Presented at 15th International Conf. on Distributed Computing Systems (ICDCS).

[2] A. DeSimone, M. C. Chuah, and O. C. Yue. (1993, Dec). Throughput Performance of Transport-Layer Protocols over Wireless LANs. Presented at Globecom.

[3] V. Jacobson. Congestion Avoidance and Control. In Proc. ACM SIGCOMM 88, August 1988.