On TCP-Induced Telehaptic Packet Loss and Jitter

Vineet Gokhale, Jan Fesl

Institute of Applied Informatics, University of South Bohemia, CZECH REPUBLIC, emails: {vgokhale, jfesl}@prf.jcu.cz

Abstract: **Telehaptic data communication (transmission of touch signals) is known to be extremely sensitive to packet loss and jitter, the primary consequences of network congestion. Existing studies have established the Quality of Service (QoS) conditions that need to be guaranteed for smooth telehaptic communication. Specifically, the telehaptic communication can tolerate no more than 10% packet loss and 10 ms jitter. In this paper, we conduct a detailed investigation of the impact of TCP cross-traffic (pre-dominant traffic on shared networks) on telehaptic packet loss and jitter. The important contribution of our study is twofold. Firstly, we discover that even during scenarios where the long term average packet loss is comfortably below its QoS limit, the instantaneous loss can far exceed this limit. Secondly, we demonstrate that the probability of jitter QoS violation increases with the number of concurrent TCP sources in the network. These effects could potentially be harmful to the telehaptic activity, thereby raising serious concerns on designing efficient communication frameworks for minimizing telehaptic packet loss and jitter on shared networks.**

Keywords: **Telehaptic communication, QoS, shared network, packet loss, jitter**

I. Introduction

Everyday activities that the humans perform are largely dependent on our sensory mechanisms that aid in learning the physical properties of any real object such as size, shape, weight, texture, hardness, smell, and so on. The touch perception forms an integral part of our sensory mechanism. When an object is held, it exerts certain forces on the hand. The muscles and the joints of the hand capture these forces and they are then transmitted to the brain, generating a perception map of the object. This sensory mechanism is the fundamental driving force behind the innumerous forms of seamless interaction between humans and the physical world. Life would be lot harder if one was to light a matchstick, drive a car, or play a game of golf without the ability to feel the physical object.

Haptics relates to the science behind the different mechanisms of perception of real objects through the sense of touch. The deep research insights in this field have led to the design of elegant electro-mechanical systems that have enabled us to interact and manipulate virtual as well as remote objects through the feeling of touch.

Telehaptic communication – the science of coding, and subsequent transmission of haptic signals over a network – has witnessed rapid progress over the past decade. Such communication has been envisaged to redefine the way we interact with a remote world. For example, a surgeon could perform telesurgery on a distant patient through a robotic telemanipulator with the feeling of touching the patient's body [1]. Telehaptic communication finds potential applications in a wide variety of other domains as well, like telemaintenance, and remote disaster management to name a few.

Figure 1 depicts a typical telehaptic communication framework over a shared network. The human operator (OP), using the force, audio, and video feedback from the remote environment, makes certain movements in an attempt to interact with and/or manipulate a remote physical object. The position and velocity signals thus generated are transmitted to the remote environment via the *forward channel*. The robotic teleoperator (TOP) at the remote location utilizes these coordinates in order to replicate OP's movements accurately. Any contact between the remote object and the TOP generates forces, which are transmitted back to the OP along with audio and video feedback on the backward channel. The presence of haptic feedback has been shown to increase the immersion into the remote environment, and further improve the precision of the telehaptic activity significantly [2].

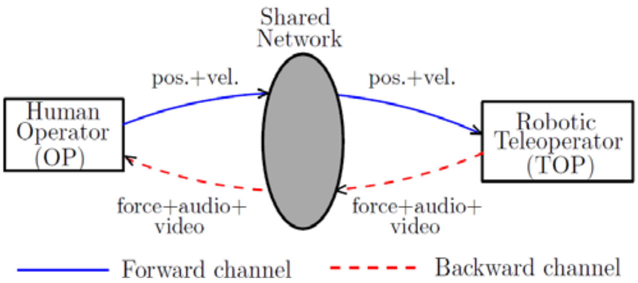

Fig. 1. Schematic representation of a telehaptic communication framework depicting the signal exchange between the human operator (OP) and the robotic teleoperator (TOP).

Naturally, such highly sensitive operations necessitate accurate replication of the OP's movements by the TOP, and also timely delivery of the feedback signals to the OP. For example, large delays in haptic feedback result in sluggish perception of the patient's body, thereby (potentially) leading to a wrong action by the surgeon. Additionally, large telehaptic jitter leads to perceiving the same remote object as having variable mass, which is absurd. Note that jitter refers to the variation in the packet delays. High packet losses may cause improper replication of the OP's movements accurately and/or OP being severely deprived of the feedback signals. Both these scenarios could have catastrophic effects on the ongoing telehaptic activity. Note that the packet loss in the network is a consequence of queue overflows during congestion. These effects can, at times, cause irreparable damage to the patient. Hence, the communication network that transfers the telehaptic feedback plays an instrumental role in determining the quality of the telehaptic interaction.

Experimental studies, such as [3], have demonstrated that the human perception of haptic feedback can tolerate a maximum packet loss and jitter of not more than 10% and 10 ms, respectively. This means that the perception of the remote environment is not hampered even if at least 90% of the telehaptic samples reach the OP/TOP with a jitter of no more than 10 ms. These telehaptic packet loss and jitter constraints that need to be satisfied for a seamless telehaptic activity are collectively known as *Quality of Service (QoS)*. For a smooth telehaptic activity, the network needs to guarantee QoS-compliance at all times. In general, QoS violations lead to deteriorated perception of the remote environment, as explained previously.

It is important to note that the work in [3] treats the packet loss as a time-average entity. In other words, the work in [3] averages packet losses over an entire telehaptic session; the authors discovered that when this long term average packet loss exceeds 10%, the users started perceiving an unacceptable deterioration in the perception of the remote environment. Note that the long term average packet loss refers to the average of the packet loss measured over the entire duration of the telehaptic session. It is worth noting that this work does not consider the characteristics of the instantaneous loss while establishing 10% as the packet loss criteria for smooth perception.

It is important to remark that in a real world scenario the perception of remote objects (potentially) depends on the instantaneous packet loss rather than the long term average loss. For example, in a few network settings the instantaneous packet loss is way higher than 10% (see Figures 4 and 5) despite its long term average value being below 10%. This means that a vast majority of the telehaptic samples (up to 80% in our simulations) do not reach the destination. As per the claim in [3], this implies that even when all packets (100%) are lost over a certain interval, the users do not feel any perceptual degradation. This is incorrect as no haptic feedback leads to improper perception of the remote world. Therefore, the instantaneous packet loss, and not the long term average loss, is a more relevant performance metric from the standpoint of perception in any telehaptic communication.

A telehaptic stream on a shared network, like the Internet, has to contend with other cross-traffic streams that are concurrently being served by the network. Hence, it is crucial to study the influence of the coexisting cross-traffic streams on the telehaptic stream in terms of instantaneous packet loss and jitter. On a shared network, the telehaptic stream is guaranteed to encounter Transmission Control Protocol (TCP) traffic since TCP amounts to over 90% of the overall traffic [4]. TCP provides a reliable data communication mode, and hence forms the cornerstone of a wide variety of Internet services that require reliable transfer of data, such as email, file transfer, web browsing, and video streaming applications like YouTube, and Netflix.

For our investigation in this paper, we consider a specific flavor of TCP named TCP NewReno [5]. TCP NewReno (or any TCP source in general) is a rate-adaptive transport layer protocol that controls its transmission rate depending on the congestion level in the network. The TCP source uses packet loss as an indicator of congestion. Based on the packet loss as detected by the source, the data rate is adapted to match the available network bandwidth, and thereby eschew underutilization the network resources. The TCP source increases its data rate until it detects a packet loss (indicating congestion). In response, it reduces the data rate in order to relieve the network, and thereby achieve congestion control. Once the source detects that the network is free, it begins to increase the data rate, and this cycle continues. As can be observed, the TCP source relies heavily on the packet loss in the network in order to learn the available network bandwidth. In fact, the working principle of TCP is itself based on inducing packet loss in the network. This behavior naturally impacts the concurrent streams in the network. In addition, the data rate variation of TCP also introduces jitter that negatively affects the telehaptic activity. In this work, we are interested in studying whether these packet loss and jitter effects of TCP have any notable impact on QoS-compliance of the telehaptic stream.

In this paper, we intend to study the impact of multiple TCP cross-traffic sources on a telehaptic stream. The objective of this investigation is to gain insights into the characteristics of the instantaneous telehaptic packet loss and jitter under the influence of coexisting TCP cross-traffic sources. The contribution of our work is as follows.

(i) We demonstrate that in a wide range of settings, even though the long term average packet loss meets the QoS criteria, the instantaneous packet loss can be much higher.

(ii) We show that the peak telehaptic jitter can far exceed the 10 ms deadline for standard network settings, and hence is extremely prone to QoS violations.

The remainder of the paper is organized as follows. In Section II, we discuss in brief a few prior works available in the literature related to the interplay between TCP and telehaptic streams. Section III describes the detailed simulation setup that we designed for our investigation. In Section IV, we present the results of our experiments, and in Section V, we state our conclusions and mention potential directions for future research.

II. Related Work

Only a handful of works have attempted to study the behavior of telehaptic streams on a shared network [6, 7, 8]. Although these works considered network cross-traffic in their performance evaluation, negligible attention is paid to the TCP streams that form a major component of the overall cross-traffic. A recent work [9] conducted a comprehensive analysis of the effects of a single TCP stream on the long term average telehaptic packet loss as well as jitter. However, this work investigates ignores the instantaneous packet loss. As explained earlier, the instantaneous packet loss forms a more important performance metric than the long term average measure. Furthermore, this analysis confines the number of concurrent TCP streams to one. Hence, the effect of multiple TCP streams on telehaptic loss and jitter remains unexplored in this work.

III. Simulation Setup

In this section, we give a detailed description of the experimental settings considered in our simulations. The goal of this section is to develop an understanding of the dynamics of interplay between TCP and telehaptic streams when the two traffic types share a single bottleneck link. We carry out our investigation using NS3 – a discrete event network simulator [10]. We use the single bottleneck network topology as shown in Figure 2. H_1 and H_2 are the OP and the TOP, respectively, of the telehaptic communication framework shown in Figure 1. [S_1, …, S_n] and [R_1, …, R_n] are the sets of n TCP sources and receivers, respectively. L_1 is the bottleneck link on the forward channel. Note that the data rate variation of TCP influenced the queue occupancy at B_1, the router at the ingress of L_1. As mentioned earlier, the TCP sources employ NewReno congestion control scheme. For telehaptic communication, we leverage the protocol proposed in [8]. It can be shown that in presence of TCP NewReno sources, a telehaptic source employing the protocol in [8] generates packets at the rate of 250 per second. The packet scheduling at the network queues is based on the standard droptail mechanism.

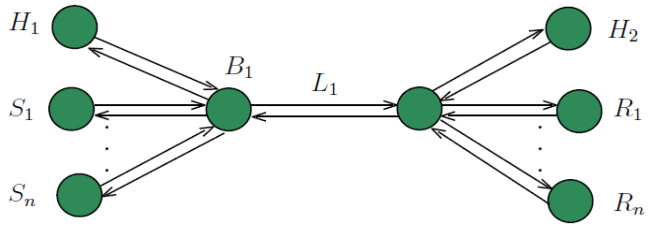

Fig. 2. Single bottleneck network topology used in our simulations. Notations: H_1 and H_2 – operator and teleoperator in telehaptic communication, respectively; [S_1 ,..., S_n] – TCP sources; [R_1 ,..., R_n] – TCP receivers; L_1 – bottleneck link; B_1 – router at the ingress of bottleneck link.

The propagation delay of each link is set to 5 ms, and hence the one-way propagation delay between a source and its corresponding receiver is 15 ms. The channel capacity of L_1 is set to 3 Mbps. The access links to L_1 have high capacities of 5 Gbps. The queue size at the B_1 is configured to 15 kB. The TCP and the telehaptic packets have sizes 578 B and 512 B, respectively, unless mentioned otherwise. For the purpose of our simulations, we consider n in the range [1, 10]. However, it is worth remarking that the observations that we make regarding the telehaptic loss and jitter hold good for higher values of n as well.

All sources begin the transmissions simultaneously at $t = 0$. We run each simulation until $t = 100$ s. Throughout the simulations, we record the packet loss and jitter encountered by the telehaptic sources.

IV. Results

In this section, we present the results of our investigation of telehaptic packet loss and jitter induced by the coexisting TCP streams. We begin by reporting the packet loss, and then move to the jitter part.

In Figure 3, we report the long term average packet loss seen by the telehaptic source by varying n over the considered range. It can be seen that the long term average packet loss is an increasing function of n. However, for $n < 10$, the long term average packet loss complies to the QoS limit of 10%. However, we note that for higher n, the average loss exceeds the QoS limit severely. For brevity, we do not report the telehaptic packet loss in the higher n regime.

We now move to studying the behavior of the instantaneous telehaptic packet loss for a specific value of n for which the average loss meets the telehaptic packet loss criteria. For this purpose, we choose $n = 10$. Note, from Figure 3, that the long term average loss for $n = 10$ is approximately 10%. From Figure 4, it can be seen that the instantaneous telehaptic packet loss varies rapidly between 0 and 50%. In addition to the peak loss measurement of 50%, it can also be seen that the packet loss QoS criteria gets violated regularly. Although we report the instantaneous packet loss only for $n = 10$, we observe similar behavior for other values of n as well. In short, even though the long term average packet loss meets the QoS criteria, the instantaneous loss can be significantly higher. This confirms our conjecture that the instantaneous packet loss should be considered as the performance metric rather than the long term average packet loss.

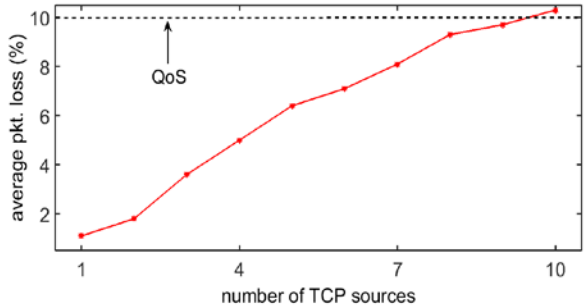

Fig. 3. Evolution of long term average telehaptic packet loss as a function of the number of TCP sources.

It is important to remark that even though the interval over which the QoS violation occurs is small (a maximum of 300 ms), this could potentially have severe artifacts considering the scale of sensitivity that a telehaptic activity, like telesurgery, requires.

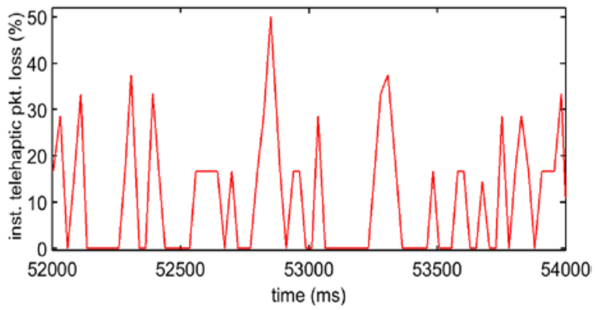

Fig. 4. Instantaneous telehaptic packet loss $n = 10$ showing significant overshoot compared to its long term average value.

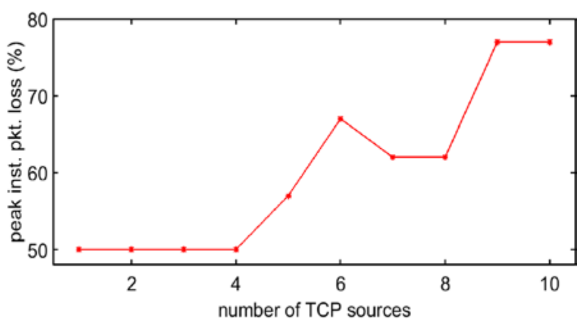

Fig. 5. Variation of peak instantaneous telehaptic packet loss as a function of the number of TCP sources in the network.

Having seen the instantaneous loss, we now turn towards determining the peak instantaneous telehaptic packet loss in the simulations. Figure 5 shows the variation of the peak packet loss in the considered range of *n*. It can indeed be observed that the instantaneous packet losses are substantially higher despite the long term average packet loss complying to the QoS requirement. Therefore, we demonstrate through experiments that any guarantees on the long term average packet loss do not imply any guarantees on the peak instantaneous packet loss. This suggests that in order to ensure a seamless telehaptic activity, one must design communication frameworks that can provide QoS guarantees on the instantaneous telehaptic packet loss.

It has been shown in the past that smaller telehaptic packets (relative to TCP packets) are less susceptible to losses [9]. Specifically, the experiments in [9] reveal that the telehaptic packets of size 137 B are rarely dropped by the network queues in presence of a single TCP source that transmits packets of size 578 B. Hence, one potential solution for mitigating the telehaptic losses is to minimize the packet sizes. The other plausible remedy could be to design priority queueing schemes that can serve packets carrying crucial telehaptic data with higher precedence over other cross-traffic streams.

We now move to the peak telehaptic jitter measurements. For the packet sizes mentioned in Section III, we notice that the peak telehaptic jitter varies in the range [2.55, 6.69] ms, which satisfies the QoS constraint. Since the jitter is known to be heavily dependent on the TCP packet size, for concreteness in exposition, we run the simulations with TCP packets of size 1042 B, which is also another standard value. In Figure 6, we plot the peak instantaneous telehaptic jitter as a function of *n*. It can be seen that the peak jitter is a non-decreasing function of *n*. Further, for $n > 5$, the jitter QoS condition is severely violated. This implies that higher the number of concurrent TCP streams, larger is the probability of violation of the telehaptic jitter QoS violation.

Based on the analysis in [9], it is reasonable to argue that the peak instantaneous haptic jitter increases with the number of concurrent TCP streams as well as their packet sizes. For QoS –compliance, one needs to theoretically determine upper bounds for these two factors. The network administrator then needs to ensure that the cross-traffic satisfies the two bounds. This guarantees satisfaction of QoS conditions for telehaptic jitter.

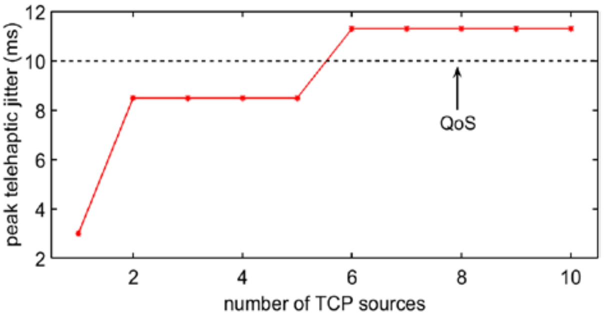

Fig. 6. Variation of peak instantaneous telehaptic jitter with TCP packets of size 1042 B.

V. Conclusions

In this paper, we conducted an extensive investigation of the interplay between TCP NewReno and telehaptic streams. We demonstrated that even though the long term average telehaptic packet loss satisfies the QoS criteria, the instantaneous loss can far exceed the QoS limit of 10%. Additionally, we showed that the telehaptic stream faces extreme jitter QoS violations for TCP packets of standard sizes. Hence, we conclude that it is crucial to monitor and control the number of TCP streams, as well as the size of TCP packets in order to achieve seamless telehaptic communication on a shared network.

In a future version of this article, we intend to propose a telehaptic communication framework that mitigates the detrimental effects of TCP sources. Also, studying the effects of other variants of TCP on telehaptic stream could be another interesting avenue for future research.

References

[1] R. Anderson, and M. Spong, "Bilateral control of teleoperators with time delay," *IEEE Transactions on Automatic Control*, vol. 34, pp. 494–501, May 1989.

[2] C. Basdogan, C. Ho, M. A. Srinivasan, and M. Slater, "An experimental study on the role of touch in shared virtual environments," in *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 7, pp. 443–460, 2000.

[3] P. Dev, D. Harris, and D. Gutierrez, A. Shah, and S. Senger, "End-to-end performance measurement of internet based medical applications," in *Proceedings of the Annual Symposium of the American Medical Informatics Association*, Nov 2002.

[4] S. Ryu, C. Rump, and C. Qiao, "Advances in internet congestion control," in *IEEE Communications Surveys and Tutorials*, vol. 5, pp. 28–39, 2003.

[5] S. Floyd, A. Gurtov, and T. Henderson, "The NewReno modification to TCP's fast recovery algorithm," 2004.

[6] R. Wirz, M. Ferre, R. Marin, J. Barrio, J. M. Claver, and J. Ortego, "Efficient transport protocol for networked haptics applications," in *Haptics: Perception, Devices and Scenarios*, Springer, 3–12, 2008.

[7] V. Gokhale, J. Nair, and S. Chaudhuri, "Opportunistic adaptive haptic sampling on forward channel in telehaptic communication," in *Haptics Symposium*, Apr 2016.

[8] V. Gokhale, J. Nair, and S. Chaudhuri, "Congestion control for network-aware telehaptic communication," in *ACM Transactions on Multimedia Computing, Communications, and Applications*, vol. 13, 2017.

[9] V. Gokhale, J. Nair, and S. Chaudhuri, "Teleoperation over a shared network: When does it work?," in *International Symposium on Haptic, Audio and Visual Environments and Games*, Oct 2017.

[10] "NS3 - The network simulator", 2011, online - http://www.nsnam.org/.