

Simplified Transport Layer Protocol for VoIP Likely Applications: Simulation and Evaluation

Omar. S. Essa

Computer Science Dep., Faculty of Science, Menoufia University, Egypt.

Abstract:

New trends in communication, in particular the deployment of multicast and real-time Audio/Video streaming applications, Internet Telephony etc. is likely to increase the percentage of non-TCP traffic in the Internet. These applications use RTP over UDP and rarely perform congestion control in a TCP-friendly manner; they do not share the available bandwidth fairly with applications built on TCP, such as Web browsers, FTP, or e-mail clients. The Internet community strongly fears that the current evolution could lead to congestion collapse and starvation of TCP traffic. For this reason, TCP-friendly protocols are being developed that behave fairly with respect to coexistent TCP flows. In this paper, we present the implementation, simulation and evaluation of a simplified transport layer protocol that we have devised. From the results obtained let me conclude that it is suitable for applications that use variable packet sizes but transmit at a constant rate (packets per second) like VoIP. Also, the protocol responds constructively to network congestion and is a TCP-Friendly. We have used ns-2 tool support for this paper [1].

Keywords: VoIP, IP, Internetworking, TCP, UDP.

1. Introduction

Most traffic in the internet uses TCP-based protocols such as Hypertext Transfer Protocol (HTTP), Simple Mail Transfer Protocol (SMTP), or File Transfer Protocol (FTP). However, the number of audio/video streaming applications such as Internet audio players, IP telephony, videoconferencing, and similar types of real-time applications is constantly growing, and it is feared that one consequence will be an increase in the percentage of non-TCP traffic. Since these applications commonly do not integrate TCP-compatible congestion control mechanisms, they treat the competing TCP flows in an unfair manner. Upon encountering congestion, all contending TCP flows reduce their data rates in an attempt to dissolve the congestion, while the non-TCP flows continue to send at their original rate. This highly unfair situation can lead to starvation of TCP traffic, or even to a congestion collapse, which describes the undesirable situation where the available bandwidth in a network is almost exclusively occupied by packets that are discarded because of congestion before they reach their destinations [2],[3].

2. TCP Friendliness

In [1], non-TCP flows are defined as TCP-friendly when “their long-term throughput does not exceed the throughput of a conformant TCP connection under the same conditions.”

TCP Friendliness for Unicast — A unicast flow is considered TCP-friendly when it does not reduce the long-term throughput of any coexistent TCP flow more than another TCP flow on the same path would under the same network conditions.

With the above definition, TCP friendliness ensures that coexisting TCP flows are not treated unfairly by non-TCP flows. Note, however, that this does not necessarily mean that all TCP and TCP-friendly flows on a bottleneck link receive the same throughput. Even competing flows that use only TCP for congestion control will often not receive the same amount of bandwidth. For example, TCP flows with different RTTs or different numbers of bottleneck nodes will transmit at different rates [4].

3. Classification Schemes for TCP-friendly behavior

3.1 Window-Based vs. Rate-Based

One possible classification criterion for TCP-friendly schemes is whether they adapt their offered network load based on a congestion window or on their transmission rate. Algorithms that belong to the window-based category use a congestion window at the sender or at the receiver(s) to ensure TCP friendliness. Similar to TCP, each transmitted packet consumes one slot in the congestion window, while each received packet or the acknowledgment of a received packet frees one slot. The sender is allowed to transmit

packets only when a free slot is available. The size of the congestion window is increased in the absence of congestion indications and decreased when congestion occurs.

Rate-based congestion control achieves TCP friendliness by dynamically adapting the transmission rate according to some network feedback mechanism that indicates congestion. It can be subdivided into simple AIMD schemes and model-based congestion control. Simple AIMD schemes mimic the behavior of TCP congestion control. This results in a rate that displays the typical short-term saw tooth-like behavior of TCP. This makes simple AIMD schemes unsuitable for continuous media streams. Model-based congestion control uses a TCP model such as the one presented in [2] instead of a TCP-like AIMD mechanism. By adapting the sending rate to the average long-term throughput of TCP, model-based congestion control can produce much smoother rate changes that are better suited to the aforementioned type of traffic. Such schemes do not mimic TCP's short-term sending rate but are still TCP-friendly over longer timescales. However, the congestion control mechanism may not resemble TCP congestion control, and great attention has to be paid to the rate adjustment mechanism to ensure fair competition with TCP or other flows [5].

3.2 End-to-End vs. Router-Supported

Many of the proposed TCP-friendly schemes are designed for best effort IP networks that do not provide any additional router mechanisms to support the protocols. Thus, they can readily be deployed in today's Internet. These schemes are called *end-to-end* congestion control. They can be further separated into *sender-based* and *receiver-based* approaches.

In sender-based approaches the sender uses information about the network congestion and adjusts the rate or window size to achieve TCP friendliness. The receivers provide only feedback, while the responsibility of adjusting the rate lies solely with the sender.

Receiver-driven congestion control is usually used together with the layered congestion control approaches. Here, the receivers decide whether to subscribe or unsubscribe from additional layers based on the congestion situation of the network. The design of congestion control protocols and particularly fair sharing of resources can be considerably facilitated by placing intelligence in the network (e.g., in routers or separate agents). Congestion control schemes that rely on additional functionality in the network are called *router-supported*. Particularly multicast protocols can benefit from additional network functionality such as feedback aggregation, hierarchical RTT measurements, management of (sub) groups of receivers, or modification of the routers' queuing strategies. *Generic router assist* (GRA) [7], for instance, is a recent initiative that proposes general mechanisms located at routers to assist transport control protocols, which would greatly ease the design and implementation of effective congestion control protocols. Furthermore, end-to-end congestion control has the disadvantage of relying on the collaboration of the end systems. Experience in the current Internet has shown that this cannot always be assumed: greedy users or applications may use non-TCP-friendly mechanisms to gain more bandwidth. As discussed by Floyd and Fall in [1], some form of congestion control should be enforced by routers in order to prevent congestion collapse. The authors present router mechanisms to identify flows that should be regulated: for instance, when a router discovers a flow which does not exhibit TCP-friendly behavior, the router might drop the packets of that flow with a higher probability than the packets of TCP-friendly flows.

While ultimately fair sharing of resources in the presence of unresponsive or non-TCP-friendly flows can only be achieved with router support, this mechanism is difficult to deploy, since changes to the Internet infrastructure take time and are costly in terms of money and effort [6].

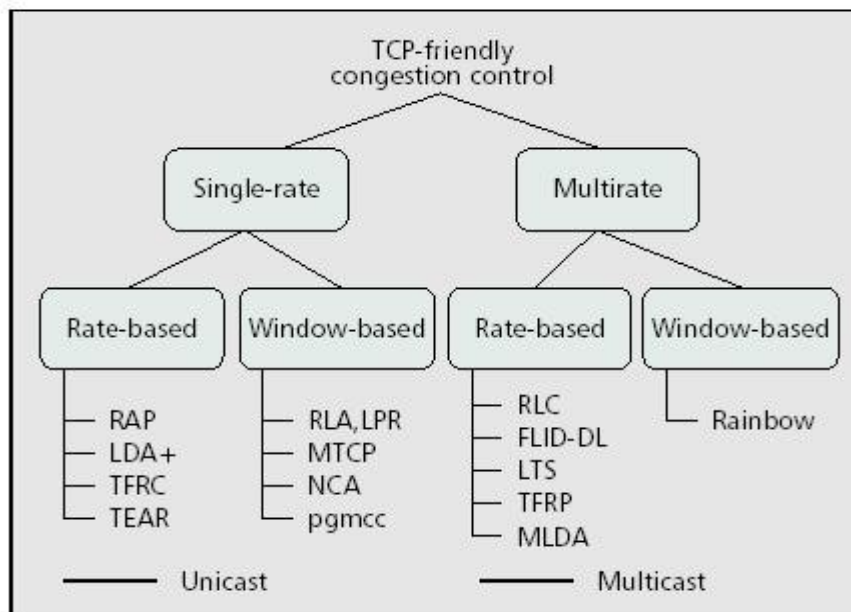


Figure 1: A classification Scheme for TCP-friendly Protocols.

4. An Overview of TFRC

TFRC is a congestion control mechanism for unicast flows operating in a best-effort Internet environment. It is reasonably fair when competing for bandwidth with TCP flows, but has a much lower variation of throughput over time compared with TCP, making it more suitable for applications such as telephony or streaming media where a relatively smooth sending rate is of importance.

The penalty of having smoother throughput than TCP while competing fairly for bandwidth is that TFRC responds slower than TCP to the changes in available bandwidth. Thus TFRC should be used only when the application has a requirement for smooth throughput, in particular, avoiding TCP's halving of the sending rate in response to a single packet drop. For applications that simply need to transfer as much data as possible in as short a time as possible; TCP is recommended, or if reliability is not required, using an Additive-Increase, Multiplicative-Decrease (AIMD) congestion control scheme with similar parameters to those used by TCP.

TFRC is designed for applications that use a fixed packet size, and vary their sending rate in packets per second in response to congestion. Some audio applications require a fixed interval of time between packets and vary their packet size instead of their packet rate in response to congestion. The congestion control mechanism proposed by TFRC cannot be used by those applications; TFRC-PS (for TFRC-Packet Size) is a variant of TFRC for applications that have a fixed sending rate but vary their packet size in response to congestion.

TFRC is a receiver-based mechanism, with the calculation of the congestion control information (i.e., the loss event rate) in the data receiver rather than in the data sender. This is well-suited to an application where the sender is a large server handling many concurrent connections, and the receiver has more memory and CPU cycles available for computation. In addition, a receiver-based mechanism is more suitable as a building block for multicast congestion control. A major advantage of TFRC is that it has a relatively stable sending rate while still providing sufficient responsiveness to compete traffic [7].

5. Problem Statement

Current TCP-friendly congestion control mechanisms such as those used in TFRC adjust the packet rate in order to adapt to network conditions and obtain a throughput not exceeding that of a TCP connection operating under the same conditions. In an environment where the bottleneck resource is a packet processing, this is the correct behavior. However, if the bottleneck resource is a bandwidth, and flows may use packets of different sizes, resource sharing then depends on packet size and is no longer fair. Now for some applications, such as Internet telephony, it is more natural to adjust the packet size, while keeping the packet rate as constant as possible.

TFRC-PS, that intends to provide a congestion control mechanism for VoIP-like applications, that vary their packet sizes and use a high constant rate of transmission, is still in its incipient stages of development and is an abstract concept.

6. Goal

To devise and implement a simplified protocol that is suitable for VOIP-like applications and at the same time responds constructively to Congestion.

7. Main Section

TFRC already implements a sophisticated mechanism for varying the rate in order to faster a constructive response to congestion. One possible approach would be to modify the mechanism slightly and scale the packet sizes instead of the rate. Such an approach does not work and here is a brief explanation.

7.1 TFRC for VoIP-like applications

For its congestion control mechanism, TFRC directly uses a throughput equation for the allowed sending rate as a function of the loss event rate and round-trip time. In order to compete fairly with TCP, TFRC uses the TCP throughput equation, which roughly describes TCP's sending rate as a function of the loss event rate, round-trip time, and packet size. We define a loss event as one or more lost or marked packets from a window of data, where a marked packet refers to a congestion indication from Explicit Congestion Notification (ECN) [8].

Generally speaking, TFRC's congestion control mechanism works as follows:

- ü The receiver measures the loss event rate and feeds this information back to the sender.
- ü The sender also uses these feedback messages to measure the round-trip time (RTT).
- ü The loss event rate and RTT are then fed into TFRC's throughput equation, giving the acceptable transmit rate.
- ü The sender then adjusts its transmit rate to match the calculated rate.

The throughput equation is:

$$X = \frac{s}{R \cdot \sqrt{2 \cdot b \cdot p / 3} + (t_RTO \cdot (3 \cdot \sqrt{3 \cdot b \cdot p / 8} \cdot p \cdot (1 + 32 \cdot p^2)))}$$

Where:

X is the transmit rate in bytes/second.

s is the packet size in bytes.

R is the round trip time in seconds.

p is the loss event rate, between 0 and 1.0, of the number of loss events as a fraction of the number of packets transmitted.

t_RTO is the TCP retransmission timeout value in seconds.

b is the number of packets acknowledged by a single TCP acknowledgement.

Different TCP equations may be substituted for this equation. The requirement is that the throughput equation be a reasonable approximation of the sending rate of TCP for conformant TCP congestion control.

The loss event rate, p , is calculated at the receiver, by aggregating the packets lost in a one RTT into single loss event. For applications that transmit at a constant rate and different packet sizes, the loss of a number of small packets in a single RTT would be aggregated into a single loss event, thus reducing the loss event rate, p . Clearly, this causes a bias towards sending small packets at a high rate. This is not fair at all since it results in lower bandwidth utilization. Hence TFRC is not suitable for Internet Telephony applications.

7.2 Packet-Size Scaling Protocol (PSP)

This is a simplified mechanism that employs a simple N-level packet scaling mechanism to respond to congestion. Basically this protocol sets a limit on the maximum packet size for an application, this limit is controlled by the present state of network congestion. The rate of transmission (Number of packets per second) remains constant and is decided before the transmission begins.

7.3 Detailed Flowchart (PSP Mechanism):

7.3.1 Sender Side: See Fig. 2.

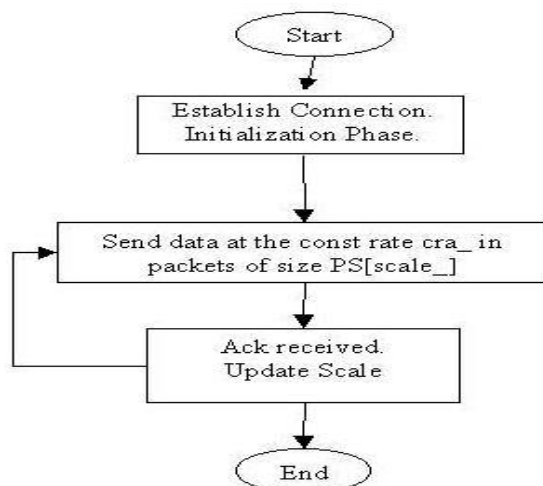


Figure 2: PSP Mechanism: Sender Side.

7.3.2 Receiver Side: See Fig. 3.

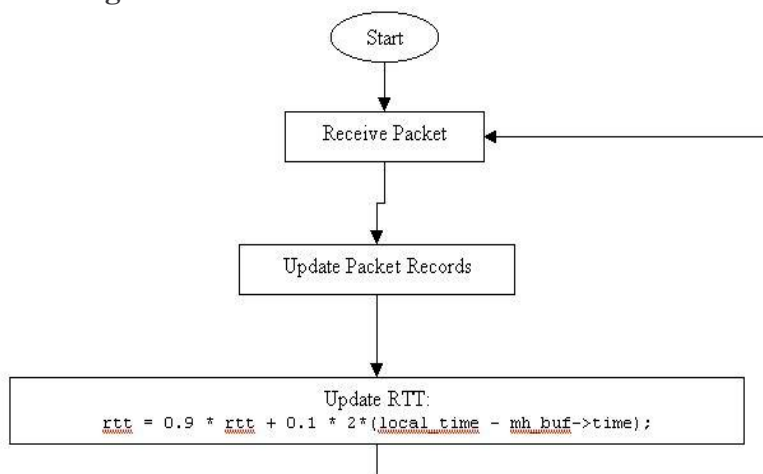


Figure 3: PSP Mechanism: Receiver Side.

7.3.3 Ack mechanism: See Fig. 4

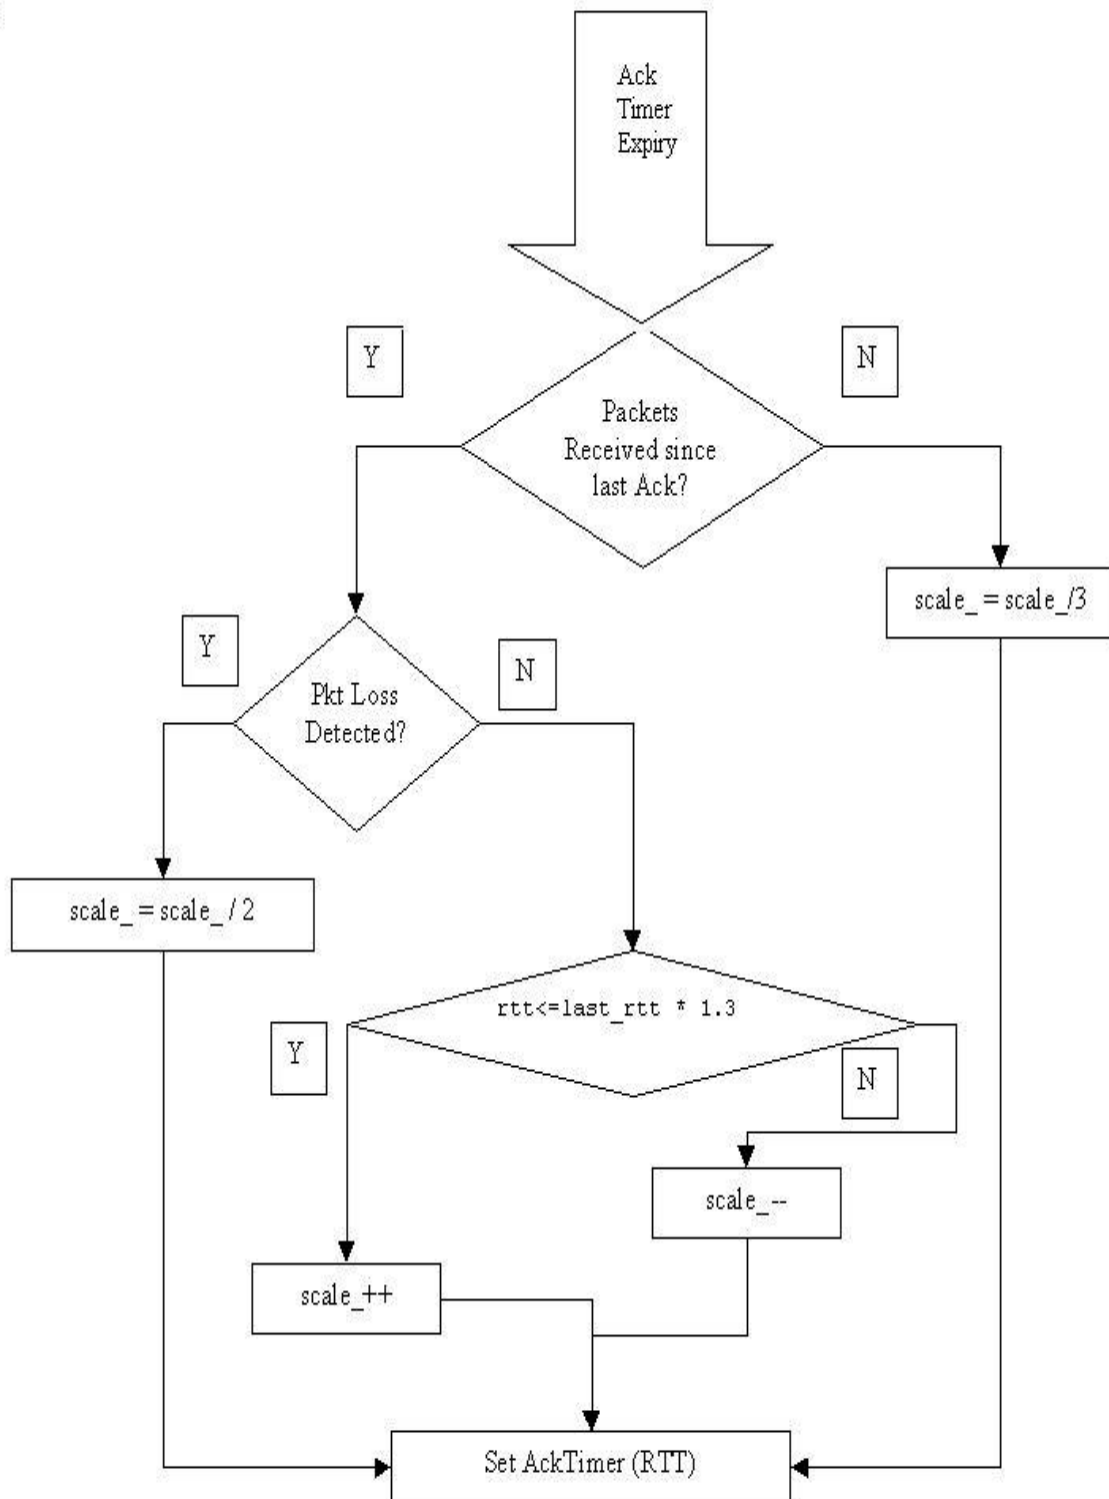


Figure 4: Ack mechanism.

7.4 PSP Protocol Sequence Diagram: See Fig. 5.

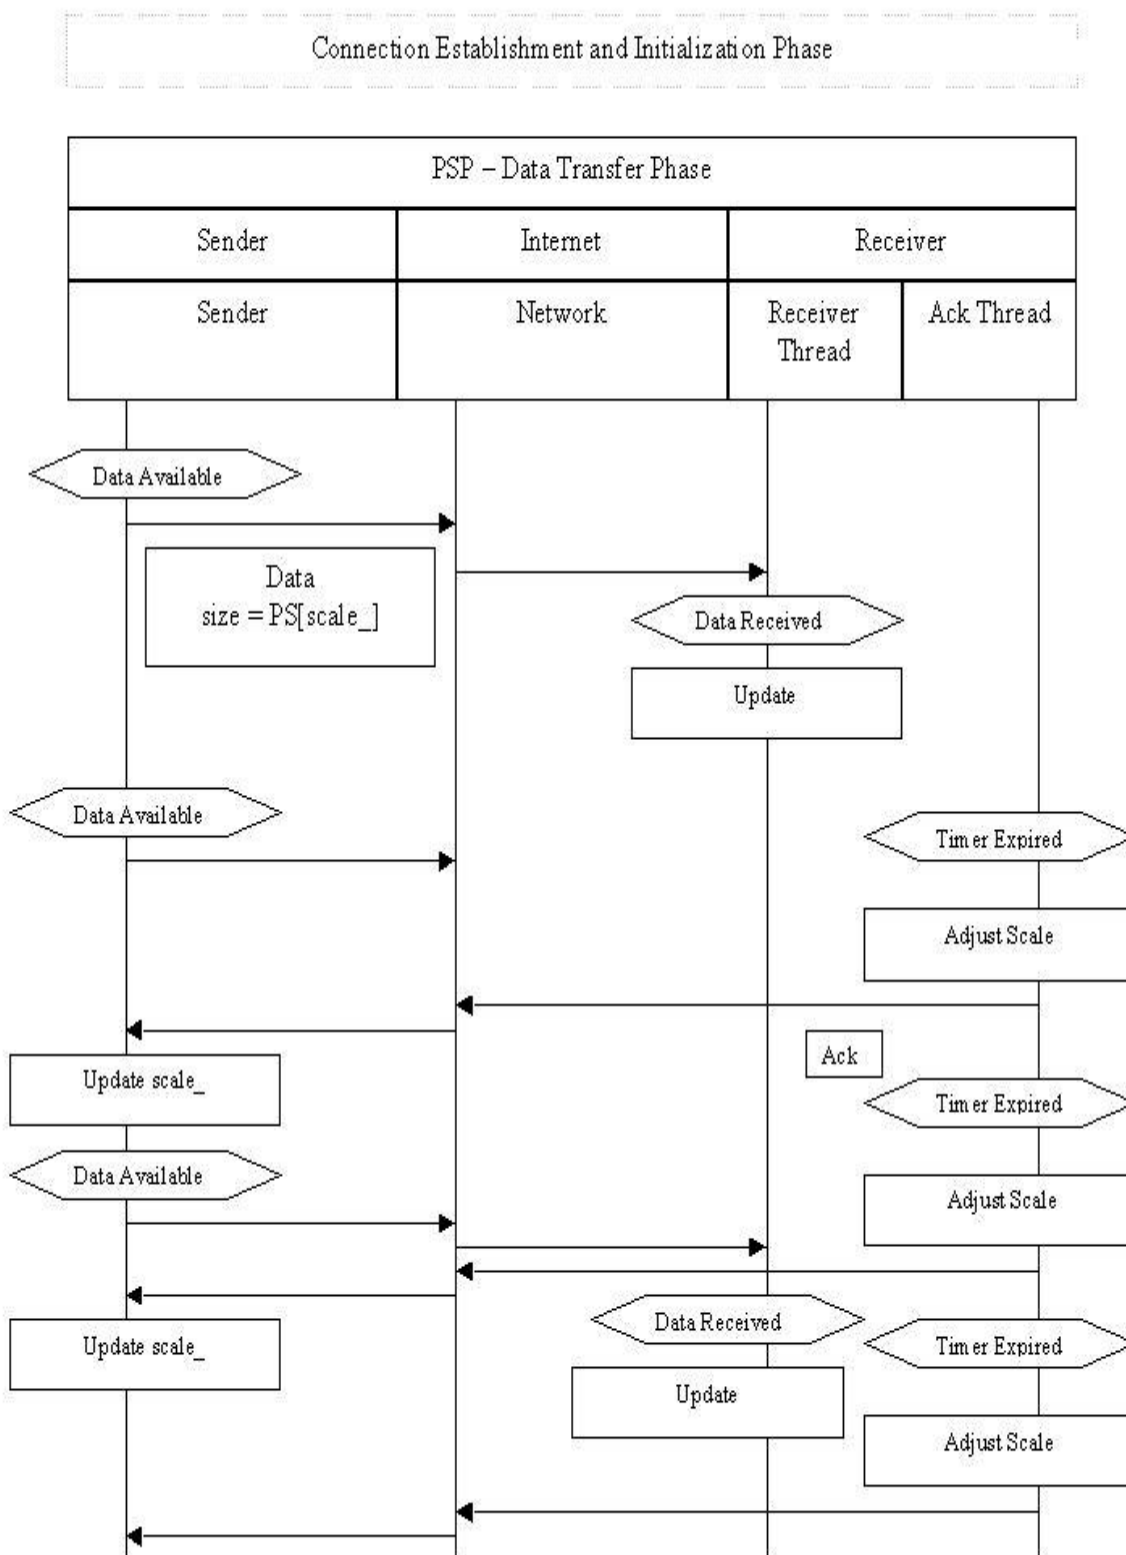


Figure 5. PSP Protocol Sequence Diagram

8. Results

8.1 TCP Vs. PSP (Throughput of the Sender)


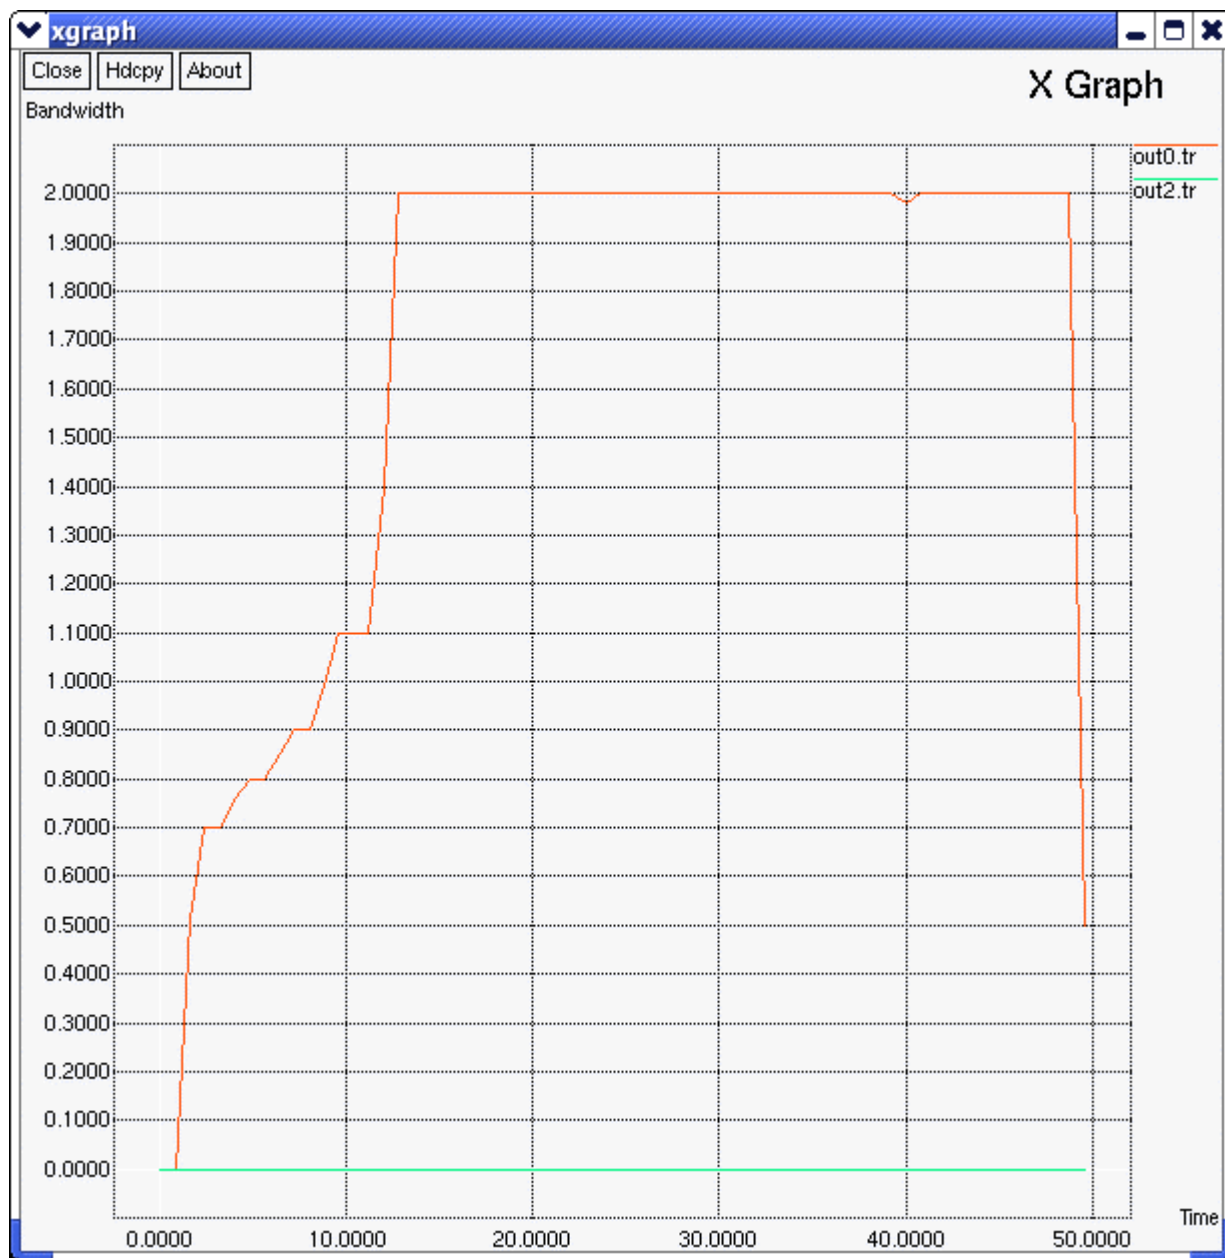


Figure 6: TCP Vs. PSP.

The graph, found in Fig. 6, lets us conclude that PSP is in fact TCP-Friendly, because the average throughput of the PSP sender is less than or equal to the average throughput of the TCP sender. Also, there is a full utilization of the available resources, because at any time the sum of the throughputs of the competing flows equals the bottleneck bandwidth.

8.2 PSP Flow (after killing the TCP-flow)



Units: Bandwidth-mb/s Time- Seconds	Legend: PSP ———— TCP ————
--	--

Figure 7: PSP Flow (after killing the TCP-flow)

The graph, found in Fig. 7, lets us conclude that the PSP-flow quickly achieves its fair share of the available resources.

8.3 Delay Variation of PSP flow (co-existing TCP-flow)

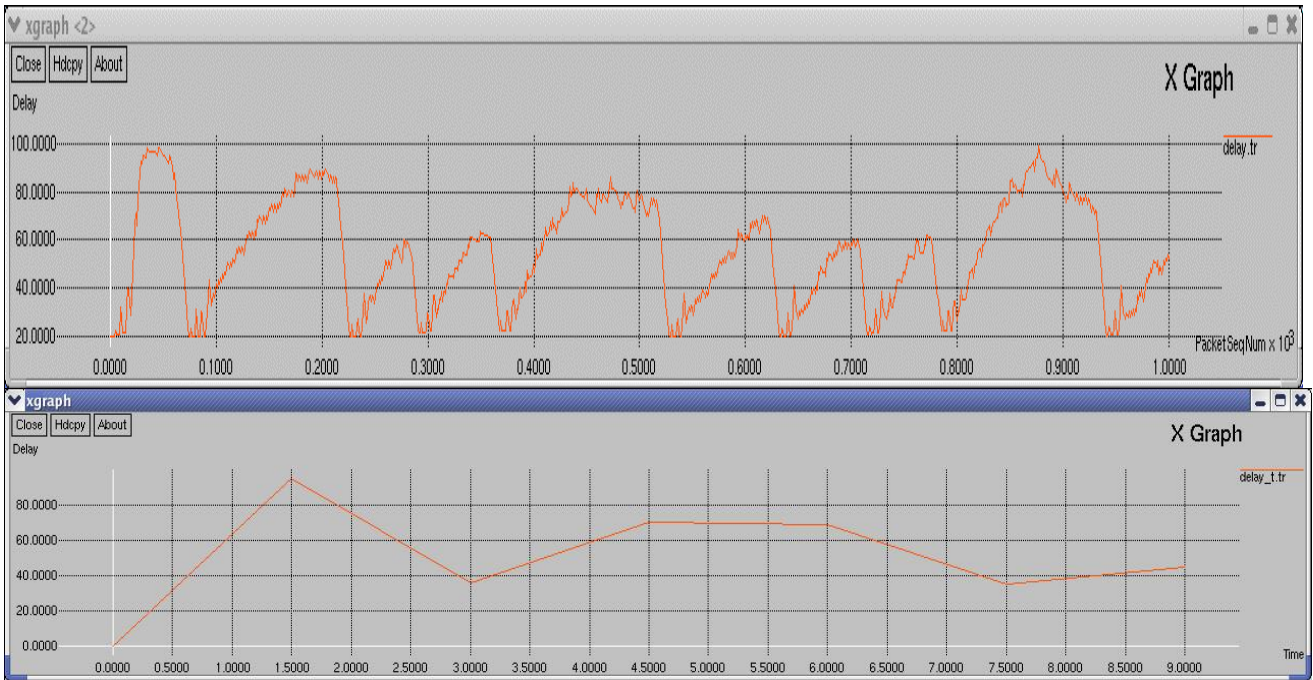


Figure 8: Delay Variation of PSP flow

Fig. 8. shows the following: the first graph plots “End-to-End Packet Delay” Vs. “Packet Sequence Number” for the first 1000 packets. The second graph plots “End-to-End Packet Delay” Vs. “Time” for 9 seconds (roughly the time in which 1000 packets are sent). Sampling interval is 1.5 seconds. These graphs let us conclude that the packet delay with PSP is suitable for VoIP-like applications, since these applications require an end-to-end delay of less than or equal to 150 ms. The delay jitter also seems to be acceptable (the second graph).

8.4 Packet Loss of PSP-flow (co-existing TCP-flow)




Figure 9: Packet Loss of PSP-flow

Fig. 9. shows the following: the first graph plots a bar graph for “Packet Loss” Vs. “Packet Sequence Number” for the first 1000 packets. A ‘1’ on the y-axis indicates that the corresponding packet was lost. The second graph plots “Packet Loss” Vs. “Time” for 9 seconds (roughly the time in which 1000 packets are sent). Sampling interval is 1.5 seconds. A 1% packet loss in the PSP-flow due to congestion was detected when it was operated along with a TCP-flow. Such a loss is acceptable for Internet Telephony applications.

8.5 At a Glance

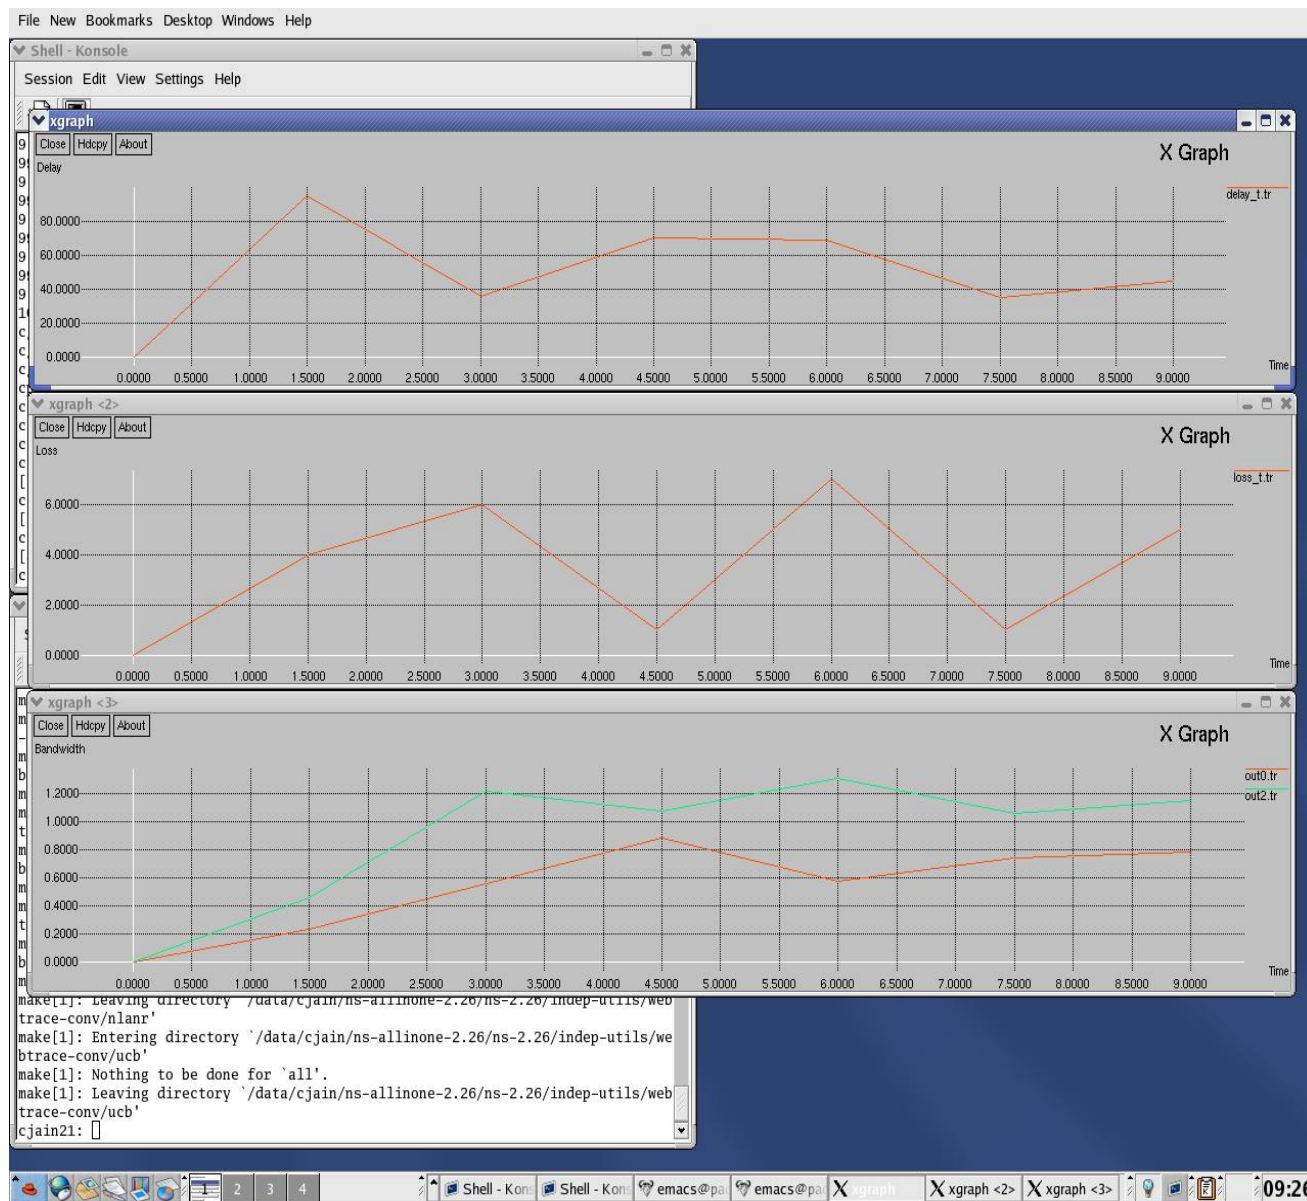


Figure 10: These graphs depict
1- “End-to-End Packet Delay” Vs “Time”.
2- “Packet Loss ” Vs “Time”.
3- “Bandwidth ” Vs “Time”.

Fig. 10. shows the following: all the graphs are for a PSP-flow that operates alongside a TCP-flow. The monitoring time is 9 seconds and the sampling interval is 1.5 seconds. A quick look at these graphs let us conclude that PSP does respond constructively to congestion. PSP perceives congestion through packet loss. As is clear from the second and the third graphs, as the packet loss increases, the transmission rate is reduced and as the congestion is cleared, the transmission rate

increases. The delay is another mechanism that we have used to perceive impending congestion and vary the rate likewise. However, this association is slightly depicted by the first and the third graphs. Please refer to “Detailed Flow-charts” for an explanation of the association between “Packet Delay” and “Transmission Rate”.

9. Discussion and Conclusions

- PSP is TCP-Friendly
- It quickly achieves its fair share of Bandwidth.
- No demands of high processing capacity at the receiver-end
- It is suitable for applications that choose to maintain a high rate at the expense of reduced packet size.
- Optimum results (for particular a simulation topology) were obtained when
 - $PS_MAX \leq MTU$
 - Constant Rate, $X \leq (\text{Bottleneck Bandwidth} / PS_MAX)$
 - $PS_MIN \geq PS_MAX / 4$
- Packet loss owing to network congestion is between 1-2%, which is good as far as VoIP applications are concerned.

10. Future Work

- A connection establishment phase, in order to automatically populate the “Preset packet-sizes” table and settle upon an optimum constant transmission rate. These parameters are important in a Bandwidth limited environment to ensure a fair share of resources.
- Study PSP behavior with RED-PD.
- The receiver side is not immune to packet re-ordering and construes that as packet loss. One way to rectify this might be to employ TCP-like mechanism. We declare a packet as lost if three packets with higher sequence numbers have already arrived.
- Add code to handle Ack lost scenario.
- What compromises does a simplified mechanism like PSP entail? What applications can do away with it?

References

- [1] McCanne, S. and Floyd, S. (McCanne, 1998) "LBNL Network Simulator,", 2003 version, 2003.
- [2] Ze. Nian Li and Mark S Drew "Fundamentals of Multimedia", Prentice-Hall Inc. Upper Saddle River, NJ, 2004.
- [3] K. RAO, Z. S. Bojkovic, and D. A. Milovanovic “Multimedia Communication Systems Techniques, Standards, and Networks”, Prentice-Hall Inc. Upper Saddle River, NJ, 2003.
- [4] Balakrishnan, H., Rahul, H., and S. Seshan, "An Integrated Congestion Management Architecture for Internet Hosts," Proc. *ACM SIGCOMM*, Cambridge, MA, September 1999.
- [5] Widmer, J.; Denda, R.; Mauve, M.; “A survey on TCP-friendly congestion control “ Network, *IEEE* , Volume:15 Issue:3, May-June 2001 Page(s): 28 –37.
- [6] S. Floyd and K. Fall, “Promoting the Use of End-to-end Congestion Control in the Internet,” *IEEE/ACM Trans. Net.*, vol. 7, no. 4, Aug. 1999, pp. 458–72.
- [7] Handley, M.; Floyd, S.; Padhye, J.; Widmer, J.; “TCP Friendly Protocol Specification (TFRC): Protocol Specification”; RFC 3448, January 2003.
- [8] Ramakrishnan, K., Floyd, S. and D. Black, "The Addition of Explicit Congestion Notification (ECN) to IP", RFC 3168, September 2001.