

OBP: Stability and Performance

Paulo Loureiro
Polytechnic Institute of Leiria
Leiria, Portugal
loureiro@estg.ipleiria.pt

Edmundo Monteiro
University of Coimbra
Coimbra, Portugal
edmundo@dei.uc.pt

Abstract

In this paper we study the stability of Open Box Transport Protocol (OBP), an explicit congestion control protocol that provides information to the end systems about the current state of the network path. To support the study we deduced the OBP transfer function to study the protocol stability, which modulates the OBP behavior at equilibrium. In the stability evaluation we use Bode and Nyquist diagrams. The OBP stability characteristics, discussed in this work, are only related with the equilibrium phase, which means when all the network capacity is in use. The results show that OBP is stable for the tested scenarios and that the OBP transfer function does not have unstable poles.

We also present an evaluation process that help to identify the OBP skills. We showed that OBP reaches high utilization of the bottleneck channel and that the OBP has fairness skills.

1. Introduction

The stability of congestion control protocols has recently been studied using fluid-flow models [7, 14, 16]. The utilization of fluid-flow models in the context of congestion control protocols showed that, for example, the TCP protocol tends to instability in fat networks [9]. The recent explicit congestion protocols have also been analyzed in terms of their stability. As an example we can refer Explicit Control Protocol (XCP) [9] and Rate Control Protocol (RCP) [5]. Recently several other TCP proposals have appeared based on the original TCP but with some modifications. These new proposals were also analyzed by the fluid-flow models. As an example, we can refer the following protocols: STCP [10], FAST [8] and HSTCP [6].

The study of stability of congestion control protocols has also been done using the known Nyquist and Bode criterions. These techniques are able to evaluate the stability of congestion control protocols and can also be used to tune some parameters of the protocols. This adjustment allows

changing the protocol behavior in regions susceptible to fall into instability. The control theory used to study the protocol's stability has nevertheless some limitations. In order to use these control tools it is necessary to represent the dynamics of the protocols by a first order system or approximately of first order. These techniques have other limitations, when applied in this area, as they are only used to study the stability when the protocol is in equilibrium. They are not used to evaluate the protocol's stability in all of its operating regions.

Open Box Transport Protocol (OBP) [11, 12] is a congestion control based on the collaboration among routers and end system. OBP estimates the transmission rate of each flow at the sender end system. Using the well known control theory we studied the stability of OBP and we shown, using the necessary approximations, that OBP is stable in the studied scenarios.

The remainder of this paper is organized as follows: Section 2 provides the related work on the study of the stability for explicit congestion control protocols; Section 3 summarizes the characteristics and the design of OBP; Section 4 explains the process of studying OBP's stability; Section 5 presents an OBP evaluation and section 6 presents the conclusions and some directions for future work.

2. Background and Related Work

The stability analysis of congestion control protocols has motivated the interest of the scientific community since many years. Several related works can be identified [7, 14, 16] about this topic. In this section we focus on the work done to study the stability of explicit congestion control protocols.

Some explicit congestion control protocols [5, 9] have been developed with the help of the traditional control techniques, which made it possible to study and to tune the stability of those protocols. The control theory of linear systems is mainly used to evaluate the stability of the explicit congestion control protocols. However, the Lyapunov control theory is also used to study the stability of explicit

congestion control protocols [4, 18, 20] when the control functions are differential, non-linear, continuous and with delay.

In [13, 19] a fluid-flow model for a general XCP network is presented, with multiple links and multiple flows. These works identify the equilibrium for the transmission rate of the XCP flows. Using the control theory, it was possible to prove that, in equilibrium, XCP avoids the filling up of the bottleneck queues. It is also demonstrated that the transmission rate defined by XCP solves the known max-min fairness criterion [19]. Finally, the control theory proves that the transmission rates defined by XCP are always the optimal.

In [3] a new method to analyze the stability of XCP and RCP congestion control protocols is presented. This work shows that the linearization process around a discontinuous equilibrium point is not a good approximation. The authors propose a new technique to study the stability, named switched linear control system with time delay. This new technique uses the Lyapunov stability theory to evaluate the behavior of the congestion control protocols. Using the previous technique, the authors of [3] prove that XCP and RCP are stable when used with the recommended configurations.

3. Open Box Transport Protocol

The OBP manages the transmission rate of the flows, supported by information about the current state of the network path. The network state is sent from the bottleneck routers to the sources. In contrast to other explicit congestion control proposals, in which the bottleneck routers define the transmission rate for all flows, the OBP updates the transmission rate in sources and individually for each flow.

For each flow, the OBP sources continuously update the transmission rate, using the knowledge of the network state and also the current transmission rate of each flow. To know the network state, the routers analyze the state of its interfaces and insert this information inside the OBP data packets, unless the packets come already marked with smaller values. More, the routers do not need to keep in memory any state information, either about flows or packets. The sources constantly receive the network state, through the reception of the ACK packets.

3.1. OBP in sources

Based on the current and continued knowledge about the network state, sources can generate packets to fill the data path, without, however, congesting the routers. Although the transmission rate is defined in the sources, the OBP also wants to reach an equal transmission rate for all the flows that have a common bottleneck. To reach its objectives the OBP has the following characteristics. The flows begin the

transmission of packets with a relatively high initial congestion window. However, the increase in the level of congestion on the network, caused by the additional packets, is quickly corrected by all the other flows. In other words, all flows are warned about the increase of traffic in the network. Then, the OBP sets for all the flows a new transmission rate based on the new state of the network. The initial congestion window is given by expression 1.

$$cwnd_{initial} = \rho * c(t) + \sigma * AB(t) \quad (1)$$

At the sources, OBP uses two different functions to calculate the transmission rate $x(t)$. Expression 2 is used if the network announces the absence of congestion ($AB(t)$ positive). Expression 3 is performed if the network announces the presence of congestion ($AB(t)$ negative). If $AB(t)$ is positive, the aggregate traffic on the network is lower than the capacity of the bottleneck link. If $AB(t)$ is negative the aggregate traffic is higher than the capacity of the bottleneck link. For each flow, the update of the transmission rate is made by the following processing.

$$\begin{aligned} & \mathbf{if}(AB(t) \geq 0) \\ x_i(t) &= xref_i(t) + xref_i(t) * \frac{AB(t-d) + \varphi * C(t)}{C(t)} * \\ & (1 - \delta * \frac{xref_i(t)}{\sqrt{(xref_i(t))^2 + (C(t) * 0.5)^2}}) \quad (2) \\ & \mathbf{else} \\ x_i(t) &= xref_i(t) - \frac{xref_i(t) * |AB(t-d)|}{C(t) + |AB(t-d)|} \quad (3) \end{aligned}$$

The constants used in these calculations take the following values: $\varphi = 0.2$ and $\delta = 0.8$. The value of the reference transmission rate is equal to the average of all transmission rates, calculated for all ACK packets received during the last RTT.

$$xref_i(t) = \sum_{perRTT} \frac{x_i(t)}{cwnd} \quad (4)$$

3.2. OBP in routers

The routers measure the state of their connections and insert this information in the OBP header of the data packets. Upon the arrival of a new packet, the routers only insert their state in the packets if it is more severe than the one previously inserted by another router of the path. Figure 1 shows the OBP header for the data packet. The network state of one interface is defined in two variables. The available bandwidth $AB(t)$ and the capacity of the interface $C(t)$. In terms of algorithm, the OBP routers divide their time in control intervals of short duration. During a control interval, all packets that enter one interface are counted and the variable $input_{traffic}(t)$ keeps this information. At the end of one control interval, the variable $input_{traffic}(t)$ contains the amount of traffic that has entered the interface.

0	7 8	15 16	23 24	31
Available Bandwidth				
Interface Bandwidth				
OBP_RTT				

Figure 1. OBP header in the data packet.

With this information and also with the information about the amount of packets that was inside the queues $queue(t)$, when the last control interval was started, routers perform the calculation presented in expression 5 for obtaining the available bandwidth in the respective interface.

$$AB(t) = C(t) - input_{traffic}(t) - \Omega * \frac{queue(t)}{avg_{rtt}(t)} \quad (5)$$

The variable $avg_{rtt}(t)$ is used to define the weight of the traffic inside the queue of the output interface. This variable is calculated in routers using the information received from the data packets (OBP_RTT).

4 Stability Analysis

In this section we deduce an expression to analyze the OBP stability. In this evaluation we use the Bode and Nyquist criterions. The OBP stability evaluation is only related with the equilibrium phase. This means, when all the network capacity is in use. To analyze the OBP stability, first, we derive the functions used by OBP to define the transmission rate of the flows, in the time domain. Next, using some approximations and simplifications, we deduce the OBP's transfer function. The transfer function is then studied to evaluate the stability. In order to use the stability models we needed to make some approximations, since those models have some restrictions in their use, as explained before. Finally, we create the Nyquist and Bode diagrams and we analyze the OBP stability for some scenarios.

Each source defines its transmission rate $x_i(t)$ whenever it receives an ACK packet. This ACK packet has the network state feedback sent by the network. The variable $ref_i(t)$ in the OBP transmission rate can be approximated per $x_i(t-d)$, in which d is the RTT.

4.1 OBP in Equilibrium

The OBP stability is studied when OBP is in equilibrium and the bottleneck capacity utilization is optimized. So, when in equilibrium, we can define the OBP transmission rate $x_i(t)$, using the following approximations:

1. Each flow that is bottlenecked in l receives the feedback $min_{l \in L} AB_l$;
2. There are N flows bottlenecked in l and each flow is using a fair share of the bottleneck capacity;

3. The available bandwidth estimated in the network changes, for each RTT, between the values in the expressions 6 and 7:

$$min_{l \in L} AB_l = 0 \quad (6)$$

$$min_{l \in L} AB_l = -C(t) * \varphi * \left(1 - \frac{\delta}{\sqrt{1+(0.5*N)^2}}\right) \quad (7)$$

Proof. Suppose that the aggregate traffic in the bottleneck link is higher than the bottleneck capacity. In this case, the sources receive the information of a negative available bandwidth. In this case the sources update the transmission rate, executing the function in expression 8.

$$x_i(t) = x_i(t-d) - \frac{x_i(t-d)*|AB(t-d)|}{C(t)+|AB(t-d)|} \quad (8)$$

Now we must prove that the aggregate transmission rate must converge to the network capacity $C(t)$. The aggregate traffic produced by the sources is given by the expression 9.

$$\sum_{i=1..N} x_i(t) = \sum_{i=1..N} x_i(t-d) - \frac{\sum_{i=1..N} x_i(t-d) * |AB(t-d)|}{C(t) + |AB(t-d)|} \quad (9)$$

This proves that if the network informs that the aggregate traffic is higher than the network capacity then, in the next RTT, the sources decrease the transmission rate and the aggregate transmission rate converges for the network capacity. If the aggregate traffic in the network is equal to $C(t)$, the network sends to the sources the information that the available bandwidth is zero $min_{l \in L} AB_l = 0$.

Now we must prove that the next aggregate traffic will be given by $min_{l \in L} AB_l = -C(t) * \varphi * \left(1 - \frac{\delta}{\sqrt{1+(0.5*N)^2}}\right)$.

Proof. If $AB(t-d) = 0$ then each source will execute the following processing.

$$x_i(t) = x_i(t-d) + x_i(t-d) * \frac{AB(t-d) + \varphi * C(t)}{C(t)} * \left(1 - \delta * \frac{x_i(t-d)}{\sqrt{(x_i(t-d))^2 + (C(t) * 0.5)^2}}\right) \quad (10)$$

After that, the aggregate transmission rate is given by the expression 11.

$$\sum_{i=1..N} x_i(t) = \sum_{i=1..N} x_i(t-d) * \left(1 + \frac{\varphi * C(t)}{C(t)} * \left(1 - \delta * \frac{x_i(t-d)}{\sqrt{(x_i(t-d))^2 + (C(t) * 0.5)^2}}\right)\right) \quad (11)$$

If $AB(t-d) = 0$, this means that in the previous RTT the aggregate traffic was equal to $\sum_{i=1..N} x_i(t-d) = c(t)$. Replacing in expression 11 we obtain $\sum x_i(t) = C(t) + C(t) * \varphi * \left(1 - \delta / \sqrt{1 + (C(t) * 0.5)^2}\right)$. In this case, the network will inform a negative available bandwidth, given by expression 7.

4. Each flow changes its transmission rate for each RTT between the two following values:

$$x_{min} = \frac{C(t)}{N} \quad (12)$$

$$x_{max} = \frac{C}{N} + \frac{C * \varphi}{N} * \left(1 - \frac{\delta}{\sqrt{1 + (0.5 * N)^2}}\right) \quad (13)$$

Proof. These results are a consequence of the point 3. If we demonstrate that, in equilibrium, the available bandwidth changes between zero and $-C(t) * \varphi * (1 - \delta / \sqrt{1 + (0.5 * N)^2})$, then this means that the transmission rate of each flow changes between the two values in expression 13 and 13. When the available bandwidth is zero the correspondent aggregate traffic is $\sum_{i=0..N} x_i(t) = C(t)$, or for each flow $x_i(t) = C/N = x_{min}$. When the available bandwidth is $-C(t) * \varphi * (1 - \delta / \sqrt{1 + (0.5 * N)^2})$ the aggregate traffic is $\sum x_i(t) = C(t) + C(t) * \varphi * (1 - \delta / \sqrt{1 + (0.5 * N)^2})$. For for each flow, the transmission rate is $x_i(t) = C(t)/N + ((C(t) * \varphi)/N) * (1 - \delta / \sqrt{1 + (0.5 * N)^2}) = x_{max}$.

From the previous observations, we can modulate the transmission rate evolution of each flow $x(t)$, with the RTT equal to d by the following expression.

$$x_i(t) = x_i(t - d) + x_i(t - d) * \varphi * \left(1 - \frac{\delta}{\sqrt{1 + (0.5 * N)^2}}\right) * \frac{4}{\pi} * \sum_{i=1..L} \left(\frac{1}{N} * \sin(N * \frac{2 * \pi}{2 * d}(t))\right)$$

In the above expression, the component $\frac{4}{\pi} * \sum_{i=1..L} \left(\frac{1}{N} * \sin(N * \frac{2 * \pi}{2 * d}(t))\right)$ is used to modulate the transitions that occur in the transmission rate.

4.1.1 Transfer Function in Open-loop

To use the Nyquist and Bode criteria the transmission rate function is converted to the Laplace domain. This transformation is given in expressions 14, 15 and 16.

$$x(t) = x(t - d) * \varphi * \left(1 - \frac{\delta}{\sqrt{1 + (0.5 * N)^2}}\right) * \frac{4}{\pi} * \sum_{i=1..L} \left(\frac{1}{N} * \sin(N * \frac{2 * \pi}{2 * d}(t))\right) \quad (14)$$

Converting.

$$s * x(s) = x(s) * e^{-ds} * \varphi * \frac{e^{ds} - 1}{s * (e^{ds} + 1)} - x(s) * e^{-ds} * \frac{\varphi * \delta}{\sqrt{1 + (0.5 * N)^2}} * \frac{e^{ds} - 1}{s * (e^{ds} + 1)} \quad (15)$$

Finally.

$$G(s) = \frac{e^{-ds}}{s} * \frac{(e^{ds} - 1)}{s * (e^{ds} + 1)} * \left(\varphi - \frac{\varphi * \delta}{\sqrt{1 + (0.5 * N)^2}}\right) \quad (16)$$

The open-loop transfer function $G(s)$ in expression 16 includes exponential terms. The Nyquist analysis of the open-loop transfer function implies that this kind of terms is not

allowed in transfer functions. To solve this situation it is necessary to use the Pade approximation [2] given by expression $e^{-ds} = \frac{(1 - sd/2 + (sd)^2/12)}{(1 + sd/2 + (sd)^2/12)}$.

The new expression for the open-loop transfer function is given by expression 17.

$$G(s) = \frac{1}{s} * \frac{12 - 6 * sd + (sd)^2}{12 + 6 * sd + (sd)^2} * \frac{(12 * sd)}{24 * s + 2 * s^3 * d^2} * \left(-\varphi + \frac{\varphi * \delta}{\sqrt{1 + (0.5 * N)^2}}\right) \quad (17)$$

Using expression 17 we can study the OBP stability, in the equilibrium state, using the Nyquist criterion and Bode criterion.

4.2 OBP: Nyquist and Bode Stability

According to the Nyquist criterion the number of unstable poles in close-loop Z is equal to the number of encirclements N around the point $-1 + 0j$ plus the number of unstable poles in open-loop P . The OBP transfer function $G(s)$ depends on the number of flows, that are sharing the bottleneck link, and it also depends on the round-trip time. For this reason we evaluate the OBP stability in some scenarios using different number of flows and round-trip times. From expression 17 it can be concluded that OBP stability is independent of the bottleneck link capacity.

4.2.1 Fixed number of flows and variable RTT

In this section we discuss the results obtained when the number of flows was 100 and the RTT varied between 0.1s and 1s. We chose to present only the maximum and minimum results of the configurations. For other configurations, the results are located between those two results. Figure 2 presents the Nyquist diagrams for the tested scenario.

Figure 2 highlights that the Nyquist line of the open-loop transfer function does not encircle the point $-1 + 0j$. More, the Nyquist line goes to the right when the RTT increases, which means, the line deviates from the instability.

The OBP stability is analyzes in Figure 3. The table includes the test conditions, the poles in open-loop and the rules of the Nyquist criterion. We can see that OBP transfer function has all open-loop poles in the left semi-plan (stable smi-plan for poles) and that all zeros and poles reduce the real and imaginary part when the RTT increases. Using the Nyquist criterion we conclude that the OBP transfer function is stable in the tested scenarios because OBP does not have any unstable poles in close-loop.

We also applied the Bode criterion. The Bode criterion defines that a system is stable if the gain margin and the phase margin are positive. In other words, the system is stable if the following two conditions are true: the gain in dB is

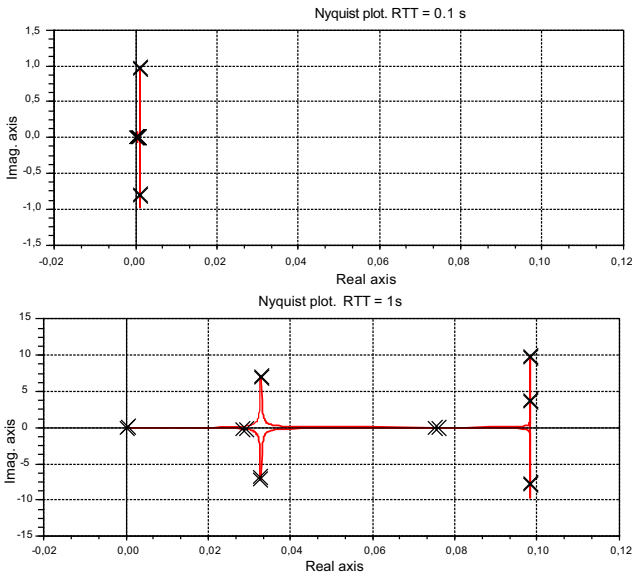


Figure 2. Nyquist diagrams.

Conditions	Open-loop poles	Z=N+P	P	N	Stable/ Unstable
1 flow, d=0.1s	0; 34.64i; -34.65i -30+17.32i; -30-17.32i	0	0	0	Stable
100 flows, d=0.1s	0; 34.64i; -34.64i -30+17.32i; -30-17.32i	0	0	0	Stable

Figure 3. OBP stability.

negative when the phase is -180° ; and the phase is greater than -180° when the gain in dB is zero. Figure 4 presents the Bode diagrams. These diagrams include the gain (Magnitude) and the Phase. When the RTT is equal to $0.1s$ and when the gain is 0 dB the phase is approximately 90° (above -180°). On the other hand, the phase line is always above -180° . These two conditions show that the system is stable for the RTT equal to $0.1s$. When the RTT is equal to $1s$ and when the gain is 0 dB the phase is approximately 70° (above -180°). On the other hand, the phase line is always above -180° . Those two conditions show that the system is stable for the RTT equal to $1s$. From the Bode analysis we can conclude that OBP is stable in the scenarios evaluated.

4.2.2 Fixed RTT and variable number of flows.

In this section we discuss the results when the RTT is $0.1s$ and the number of flows varies between 1 and 100. As explained before, we only present the maximum and minimum results of the configurations. Figure 5 presents the Nyquist diagrams for the scenario evaluated. Figure 5 highlights that the Nyquist line of the OBP open-loop transfer




Figure 4. Bode diagrams.

function does not encircle the point $-1 + 0j$. When the number of flows is increased the Nyquist line moves to the right and does not encircle the point $-1 + 0j$. The stability analysis is included in Figure 6. It can be seen that the OBP transfer function has all open-loop poles in the left semi-plan (stable semi-plan). The localization of the poles and zeros does not change when the number of flows is changed. Using the Nyquist criterion we conclude that the OBP transfer function is stable in the tested scenarios because OBP does not have any unstable poles in close-loop.

Figure 7 presents the Bode diagrams including the gain (Magnitude) and the phase for the two scenarios identified above. When the number of flows is equal to 1 and the gain is 0 dB the phase is approximately -90° (above -180°). On the other hand, when the phase line is -180° the gain is -60 dB . Under these conditions the system is stable. When the number of flows is equal to 100 and the gain is 0 dB the phase is approximately 90° (above -180°). On the other hand, the phase line is always above -180° which indicates that the system is stable. From the Bode criterion we can conclude that OBP is stable in the evaluated scenarios.

5 Evaluation

The OBP was evaluated in the network simulator NS-2 (Version 2.31) [17] with the OBP implementation. The evaluation consisted of a group of tests that helped to identify the OBP skills. The test scenarios used the one bottle-




Figure 5. Nyquist diagrams.

Conditions	Open-loop poles	Z=N+P	P	N	Stable/ Unstable
1 flow, d=0.1s	0; 34.64i; -34.65i -30+17.32i; -30-17.32i	0	0	0	Stable
100 flows, d=0.1s	0; 34.64i; -34.64i -30+17.32i; -30-17.32i	0	0	0	Stable

Figure 6. OBP stability.

neck topology. The OBP performance was also compared with other congestion control protocols, TCP Reno [1], TCP SACK [15] and XCP [9].

5.1 Link Utilization

This section presents three groups of tests: varying the bottleneck link capacity; varying the number of FTP flows; and varying the propagation delay. The results are relative to the parameter "use of the bottleneck link". The results must be interpreted in the following way. Each point in the figures corresponds to the result of one test. This means that, for example, in Figure 8, each protocol was evaluated by 19 independent tests and this figure includes 76 tests for the four protocols. More, the images below have the results of four protocols. Every test was executed alone. The results of the four protocols are together in the same images because this situation helps the comparative interpretation. The duration of every test was 100 s.

In the group varying the bottleneck link capacity, the bottleneck capacity varied from 1 Mb/s to 1 Gb/s. The other configurations used the following values: the round-




Figure 7. Bode diagram.

trip propagation delay was 100 ms; 20 FTP flows in the forward path and 20 more in the reverse; 100 new web-based flows per second; 10 voice flows.

Figure 8 presents the result in terms of the bottleneck link utilization. Protocols Reno and SACK did not use all the available bandwidth. This trend is emphasized after the 20 Mb/s. The OBP reached levels of occupation always above 90%. This happened with both low and high bandwidth.

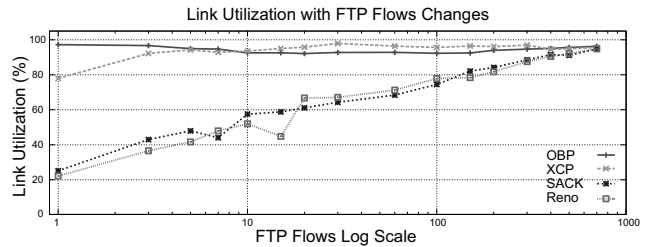


Figure 8. Link utilization, bandwidth changes.

In the group varying the number of FTP flows, we intend to evaluate the OBP when we vary the number of long FTP flows. The OBP traffic (FTP) varied between 1 and 700 flows, in the forward direction. In the reverse direction, the tests used 100 flows. The remaining configurations used the following values: the RTT was 100 ms; 100 new web-based flows per second; 50 voice flows and 50 streams of video in the forward direction and 50 in the reverse direction.

Figure 9 shows the utilization of the bottleneck channel. With few FTP flows, the Reno and SACK cannot use all the available bandwidth. The OBP nearly used all the network capacity, with few flows or with many FTP flows. In conclusion, the OBP performance is independent of the number of FTP flows that share the network. The XCP results are equivalent to the OBP.




Figure 9. Link utilization, FTP flows changes.

In the group varying the propagation delay, the tests were performed to evaluate the OBP and other protocols when the propagation delay varies. These tests used the following settings: the capacity of the bottleneck link was fixed in 1 Gb/s; 100 FTP flows in the forward direction and 100 in the reverse direction; 100 new web-based flows per second; 50 flows of voice; 50 flows of streaming video in the forward direction and 50 in the reverse. The RTT varied between 8 ms and 700 ms.

Figure 10 shows the occupation of the bottleneck channel obtained by each protocol. For the configurations with long propagation delays, the TCP Reno and SACK had difficulty in using all the available bandwidth. For all the configurations, the OBP reached high levels of occupation of the bottleneck link, above 90%. The XCP reached equivalent results as the OBP.

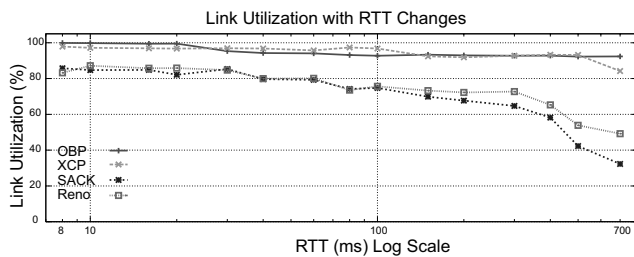


Figure 10. Link utilization, RTT changes.

5.2 Fairness: flows with different RTTs

In this section we wanted to know how the OBP manages the flows with different RTTs. For such, we generated 10 flows, all with different RTTs, from 50 ms up to 95 ms.

The capacity of the bottleneck link was 100 Mb/s. Figure 11 shows the transmission rate achieved by each flow. We can verify that each flow approximately used 10 Mb/s and this means that the transmission rate achieved by each flow is independent of the RTT. As a result, we can say that the OBP is fair for the flows with short or long delays.

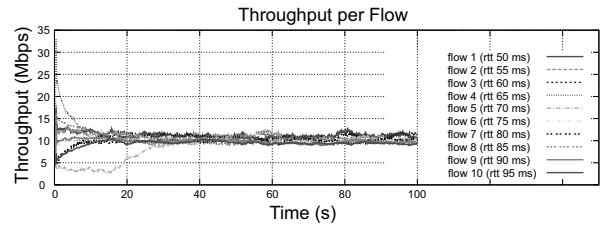


Figure 11. Throughput per flow.

Figure 12 shows the size of the congestion window of each flow. We can confirm that the flow with the highest RTT (95 ms) has the greatest congestion window. Looking at the extremes, the flow 1 (RTT = 50ms) has a window of approximately 65 packets. The flow 10 (RTT= 95 ms) has a window of approximately 120 packets. For all the flows, if we consider the value of the RTT and the value of the congestion window we conclude that all the flows obtained an equal percentage of the bandwidth. From these results we can confirm that the OBP ensures fairness to the flows with different RTTs.




Figure 12. Congestion window per flow.

5.2.1 Convergence

For evaluating the convergence, we initiated the flows with a separation of 200 seconds. With this approach, we wanted to know how the OBP handles the new flows in presence of flows already stabilized. For these tests, we used the dumb-dell topology, with 100 ms of propagation delay.

The results show that, every time a new flow is initiated, the already stabilized flows reduce their transmission rates. This behavior is very important and enables the new flows to increase the transmission rate. After some time, all the flows reach similar transmission rates. At this time and in balance, all the flows remain stabilized and share the network.

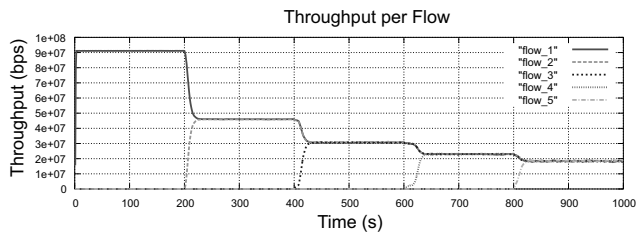


Figure 13. Throughput per flow.

6 Conclusion and Future Work

In this paper we studied the stability of Open Box Transport Protocol (OBP). OBP is an explicit congestion control protocol that provides information to the end systems about the current state of the network path. OBP estimates the flow transmission rate, at the sender end systems, using the current state of the network path.

The OBP transfer function was derived to study the stability, which modulates the OBP behavior in equilibrium. For the stability evaluation we used the Bode and Nyquist criterions. The results shown that OBP is stable for the tested scenarios and that the transfer function do not have unstable poles.

We also present an evaluation process that help to identify the OBP skills. We showed that OBP reaches high utilization of the bottleneck channel and that the OBP has fairness skills. The results showed that the OBP performance outperforms TCP Reno and SACK and XCP in some conditions. As part of our future work, we plan to implement the OBP protocol in the linux operating system.

References

- [1] M. Allman, V. Paxson, and W. Stevens. Tcp congestion control. *RFC 2581*, 1999.
- [2] G. Baker and P. Graves-Morris. Pad approximants. *Cambridge U.P., Cambridge (GB), second edition*, 1996.
- [3] H. Balakrishnan, N. Dukkipati, N. McKeown, and C. Tomlin. Stability analysis of explicit congestion control protocols. *Communications Letters, IEEE*, October 2007.
- [4] S. Deb and R. Srikant. Global stability of congestion controllers for the internet. *IEEE Transactions on Automatic Control*, June 2003.
- [5] N. Dukkipati, M. Kobayashi, R. Zhang-Shen, and N. McKeown. Processor sharing flows in the internet. *In Thirteenth International Workshop on Quality of Service*, June 2005.
- [6] S. Floyd. Highspeed tcp for large congestion windows. *RFC 3649*, December 2003.
- [7] C. Hollot, V. Misra, D. Towsley, and W.-B. Gong. A control theoretic analysis of red. *In Proceedings of IEEE INFOCOM 2001*, April 2001.
- [8] C. Jin, D. Wei, and S. Low. Fast tcp: motivation, architecture, algorithms, performance. *In Proceedings of IEEE INFOCOM 2004*, March 2004.

- [9] D. Katabi, M. Handley, and C. Rohrs. Congestion control for high bandwidth-delay product networks. *In Proceedings of ACM SIGCOMM 2002*, August 2002.
- [10] T. Kelly. Scalable tcp: Improving performance in high-speed wide area networks. *Computer Communication Review*, April 2003.
- [11] P. Loureiro, S. Mascolo, and E. Monteiro. Open box protocol (obp). *In proceedings of the High Performance Computing and Communications (HPCC) 2007*, September 2007.
- [12] P. Loureiro and E. Monteiro. What if the end systems knew the bandwidth available in the network? *In proceedings of the IEEE Global Communications Conference (IEEE GLOBECOM 2008)*, November 2008.
- [13] S. Low, L. Andrew, and B. Wydrowski. Understanding xcp: equilibrium and fairness. *In Proceedings of INFOCOM 2005*, 2005.
- [14] S. Low, F. Paganini, J. W. S. Adlakha, and J. Doyle. Dynamics of tcp/red and a scalable control. *In Proceedings of IEEE INFOCOM 2002*, June 2002.
- [15] M. Mathis, J. Mahdavi, S. Floyd, and A. Romano. Tcp selective acknowledgement options. *RFC 2018*, 1996.
- [16] V. Misra, W.-B. Gong, and D. Towsley. A fluid-based analysis in a network of aqm routers supporting tcp flows with an application to red. *In Proceedings of ACM SIGCOMM 2000*, September 2000.
- [17] NS-2. Network simulator, www.isi.edu/nsnam/ns/.
- [18] R. Srikant. The mathematics of internet congestion control. *ISBN: 978-0-8176-3227-4*, 2004.
- [19] P. Wang and D. Mills. Simple analysis of xcp equilibrium performance. *Information Sciences and Systems*, March 2006.
- [20] L. Ying, G. Dullerud, and R. Srikant. Global stability of internet congestion controllers with heterogeneous delays. *In Proceedings of the American Control Conference 2004*, June 2004.