DESCRIPTION FOR THE GENERAL PUBLIC

Routers and switches that make up computer networks and the Internet have inside buffers for packets (packets are elementary chunks of data traveling through the network). These buffers are simply spaces in electronic memory. In a buffer, packets can wait for some time before being sent through an output link to the next router/switch or to the final end user.

The buffers in networking devices play an important role. It can happen that more packets arrive at a router in a short time than the router can send quickly through the output link. If there were no buffer, many packets from such a burst would have been deleted, causing problems. If there is a buffer, they do not have to be deleted – they can wait in the buffer before being transmitted further.

Having buffers is necessary but not completely cost-free. A packet entering a router is not sent immediately through the output link but joins the queue of other packets in the buffer and waits patiently for its turn for transmission. It takes some time before it moves to the head of the queue and can be transmitted. Therefore, the time spent in the buffer (queueing delay) adds to the delivery time of the packet, from the sender to the receiver. Furthermore, there may be many routers on the delivery path of a packet, so many such queueing delays are added. Finally, the buffers do not eliminate packet losses (deletions) completely, just reduce them significantly. Every buffer may get full occasionally, and a packet arriving when the buffer is full is lost

In this way, the buffers have an important impact on the performance of the network. This impact has been investigated for many years by many researchers. In particular, the queueing delay, queue size, packet losses, and other characteristics were studied, depending on the buffer size and packet traffic arriving at the buffer.

What is important is that the whole aggregated traffic was considered in these studies. In such aggregated traffic, there is a mix of flows of packets from many individual users. Each flow may be composed of packets from a teleconference, an online game, a Netflix movie, an email, or a web page. As we can guess, each flow can have different properties – e.g., packet sizes and times between consecutive packets can be very different among flows.

The mentioned previous studies enabled us to understand very well how buffers influence the performance of the network as a whole, on an averaging basis, i.e., taking into account the aggregated traffic.

In this project, we want to understand how buffers influence every individual flow passing through the buffer, depending on the character of this flow. It is a surprising fact, confirmed in simulations, that passing through the same buffer, two different flows in the aggregated traffic may be affected in very different ways. For instance, the packet loss ratio in one flow can be 1%, and in another flow, it can be 5%. This means that some users may perceive much better network performance than the average, due to the particular character of their flows. Other users may perceive much worse performance than the average

To understand that, we want to find the values of the most important performance characteristics (queue size, queueing delay, throughput, loss ratio) for individual flows, not the whole traffic. For instance, we want to find the average number of packets present in the buffer, belonging to each individual flow. Then, we want to check how these characteristics depend on the properties of the flows and parameters of the buffering mechanism. Realizing these objectives will enable conscious parameterization of buffering mechanisms, such that they would have a predictable impact on flows of a particular type.

Finally, a simple shared buffer, at which the packets are deleted only when the buffer is full, is not the only possible design of the buffering mechanism. Other mechanisms, exploiting many buffers, or using different packet deletion policies, are considered by researchers and engineers. In the project, we will analyze four such mechanisms, of different types, in every case dealing with the performance of the mechanism with respect to particular flows.

To obtain the performance characteristics of each flow of buffering mechanisms, three different methods will be exploited: mathematical modeling, computer simulations, and measurements of real traffic in a networking lab.

The project results will be published in renowned scientific journals and presented at international conferences dedicated to computer networks.