DESCRIPTION FOR THE GENERAL PUBLIC

From the very beginning of the Internet it has been assumed, that the network does not provide any guarantees about the delivery of data packets to end users. In other words, the network tries to deliver the packets, but it is not sure, how long will it take, and even not sure, that all the packets will be eventually delivered. In fact, packet losses happen all the time in the Internet. A significant fraction of packets are deleted at network nodes (routers).

Usually, end users do not suffer from these losses. This is due to the fact, that the Internet has built-in, in the TCP protocol, mechanisms for retransmitting the lost packets. Namely, every lost packet can be retransmitted, even a few times, if necessary. For this reason, the emails we read and the web pages we browse are usually complete - there are no gaps nor glitches in them.

However, the problem with losses gets a bit worse, when there is no time to wait for the retransmitted packet. Imagine, for instance, an Internet voice call. If a packet with the coded voice is lost, its retransmission may take too much time to maintain a nice, fluent conversation.

But even in such a voice call we can tolerate quite well some losses, if they occur separately. The problem becomes really severe when the losses occur in groups, one after another. This can be demonstrated on the following figurative example. Let us assume that the following message is transmitted in the network (for simplicity, one character is equivalent to one packet):

THE_BIKE_IS_BLUE

Assume first that four separate characters are lost in this transmission:

TE_BIK_S_BLU

As we can see, it is still possible to guess correctly the meaning of the message, even though the loss ratio is pretty high (25%). Now assume the second scenario, in which the same number of characters are lost, but in two groups of two characters:

THE_KE_IS_UE

It is very hard to guess correctly the meaning of the message.

More or less the same happens in real-time voice and video transmissions. Losing one video frame out of 25 in every second may be even unnoticed. On the other hand, 20 frames lost in a row will be definitely noticed, as an unpleasant gap in the video.

For all these reasons, in this project we will study the parameter that characterizes the tendency of packet losses to cluster together. This parameter is called the burst ratio. In particular, we will try to find, what are the values of the burst ratio in the Internet, which factors make it high and which do not, and what is its influence on the quality of real-time multimedia transmissions. As the bare burst ratio does not provide the full information about the losses, some other statistics of the loss process, which reflect its more subtle properties, will be studied as well.

Packets are lost in Internet routers for a well-known reason. Namely, the incoming packets are stored in a buffer before transmission via the router's output interface. If this buffer is full, the arriving packets are deleted. Therefore, the studies of the burst ratio have to be carried out using mathematical models of this mechanism. Properly parameterized and tuned, such models can mimic precisely the actual reason of losses - buffer overflow events. The main complication in this analysis is the fact, that Internet traffic has a very complex statistical structure, in which several strange phenomenons occur, e.g. a long-range dependence of the intensity. Fortunately, scientists developed some useful mathematical models of such traffic, which can be used as a component of the whole model of packet loss at the router's buffer.

After solving the mathematical models and writing programs for numerical calculations, we will be able to predict, what values of the burst ratio can be expected in the Internet, assuming realistic parameterizations of traffic and queueing mechanisms. Naturally, the theoretical studies will be accompanied by actual measurements of the burst ratio in Internet nodes available to us. The measured values will be then compared with the theoretical predictions.

But the project is devoted not only to obtain a precise knowledge of the value of the burst ratio and the factors that can potentially make it high. We will also study a method of decreasing the burst ratio in the Internet, based on the application of the dropping function. Namely, a packet arriving to the queue at the network node can be dropped (deleted) with some probability, even if the buffer is not full yet. It is already known that such policy may have several positive effects on Internet transmissions. It is likely, that this mechanism will serve well yet another important purpose - decreasing the burst ratio - but it has to be properly tuned and parameterized first.