WestwoodTCP

TCP Westwood

The current Standard TCP implementation rely on packet loss as an indicator of network congestion. The problem in TCP is that it does not possess the capability to distinguish congestion loss from loss invoked by noisy links. As a consequence, Standard TCP reacts with a drastic reduction of the congestion window. In wireless connections overlapping radio channels, signal attenuation, additional noises have a huge impact on such losses.

TCP Westwood (TCPW) is a small modification of Standard TCP congestion control algorithm. When the sender perceive that congestion has appeared, the sender uses the estimated available bandwidth to set the congestion window and the slow start threshold sizes. TCP Westwood avoids huge reduction of these values and ensure both faster recovery and effective congestion avoidance. It does not require any support from lower and higher layers and does not need any explicit congestion feedback from the network.

Measurement of bandwidth in TCP Westwood lean on simple relationship between amount of data sent by the sender to the time of receiving an acknowledgment. The estimation bandwidth process is a bit similar to the one used for RTT estimation in Standard TCP with additional improvements. When there is an absence of ACKs because no packets were sent, the estimated value goes to zero. The formulas of estimation bandwidth process is shown in many documents about TCP Westwood. Such documents are listed on the UCLA Computer Science Department website.

There are some issues that can potentially disrupt the bandwidth estimation process such a delayed or cumulative ACKs indicating wrong order of received segments. So the source must keep track of the number of duplicated ACKs and should be able to detect delayed ACKs and act accordingly.

As it was said before the general idea is to use the bandwidth estimate value to set the congestion window and the slow start threshold after a congestion episode. Below there is an algorithm of setting these values.

```
if (n duplicate ACKs are received) {
   ssthresh = (BWE * RTTmin)/seg_size;
   if (cwin > ssthresh) cwin = ssthresh; /* congestion avoid. */
}
```

where

- `BWE` is a estimated value of bandwidth,
- `RTTmin` is the smallest RTT value observed during the connection time,
- `seg_size` is a length of TCPs segment payload.

Additional benefits

Based on the results of testing TCP Westwood we can see that its fairness is at least as good, than in widely used Standard TCP even if two flows have different round trip times.
The main advantage of new implementation of TCP is that it does not starve the others variants of TCP. So, we can claim that TCP Westwood is also friendly.

References

- TCP Westwood proposed by UCLA Computer Science Department website

– Main.WojtekSronek - 27 Jul 2005