Digital Rapids Copper vs TCP

How FTP (TCP/IP based) data transmission works

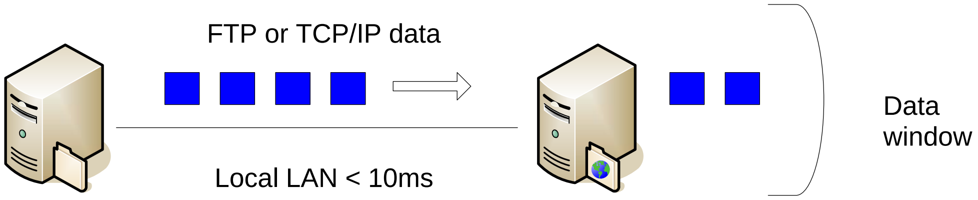

TCP sends data serially, one packet at a time, in order, until the receiver has a full buffer (receive window).

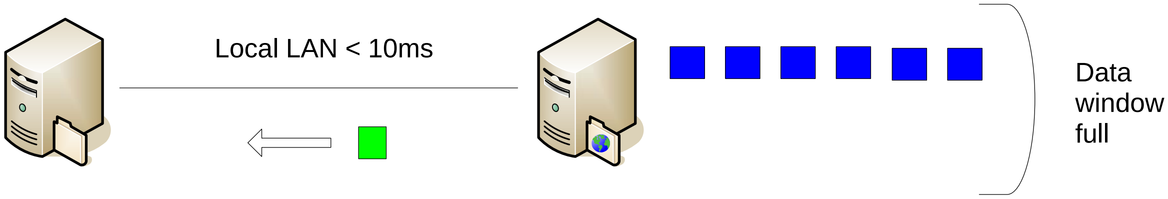

When the receive window is full, the receiver sends an acknowledge (Ack) back to the sender and the process can start over. The time it takes each packet to arrive at the receiver or the time it takes the acknowledge to return to the server is called latency. The latency of a typical local area network or LAN is less than 10ms.

How transmission rate is affected by latency

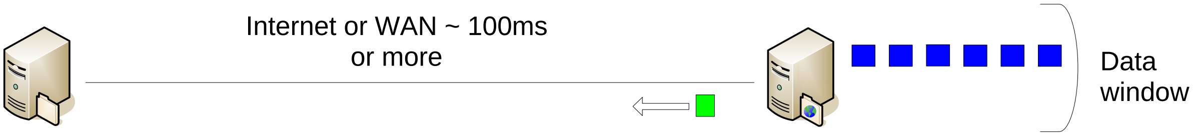

In a wide area network or WAN (such as the public internet), the latencies are much higher, so each packet takes longer to arrive as does each acknowledge. This drastically slows down the overall throughput of the network. A good resource for latency through public networks is the Internet Traffic Report (http://www.internettrafficreport.com). This site can show you the global or regional average network response time (latency). At the time of this document writing, latency was typically over 130ms in North America. Copper has a speed advantage over typical network protocols when the latency gets above 50ms. At 100ms it can be many times faster than the typical ftp protocol.

How transmission rate is affected by loss

Another factor in determining the overall potential throughput of any network is loss. Loss can be caused by many factors but most typically congestion. Congestion is the overloading of any single point on the network between the server and the receiver. It can happen locally, to either end, or anywhere in between.

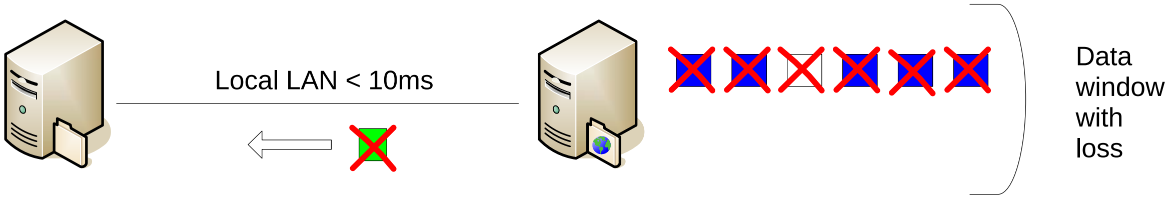

Because TCP/IP requires that all packets arrive in order, any single lost packet means that the entire receive buffer must be re-transmitted. This can drastically affect overall network performance. The way that TCP/IP deals with loss is by lowering the utilization (transmission rate) of the network, further reducing performance. Sporadic, unpredictable loss will result in the network performance spiking and falling, reducing throughput dramatically. High amounts of loss can cause TCP/IP based applications (i.e. FTP) to completely fail.

The Internet Traffic Report (http://www.internettrafficreport.com) is again a good site to use to find out the average packet loss in any given region. Typically loss in North America averages between 2 and 3%. In Europe the average loss is similar. In Asia the average packet loss can be as high as 10%.

How transmission rate is affected by loss and latency combined

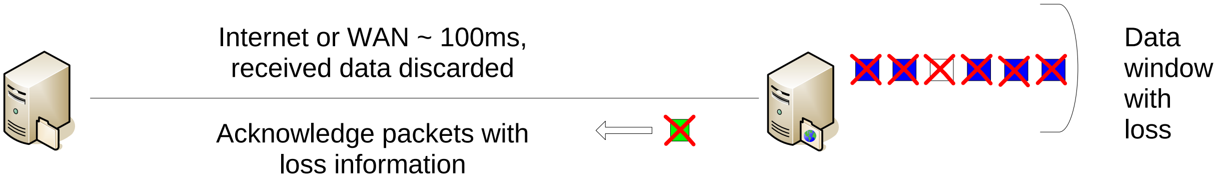
Internet or WAN ~ 100ms, received data discarded
Acknowledge packets with loss information
Data window with loss

Every network is susceptible to loss, even local area networks. The more geographically separated two networks are, the higher the latency will be. Loss and latency together combine to absolutely destroy TCP/IP performance. Typically it is very difficult to get more than 20% utilization of bandwidth over a link with high latency alone. If the loss becomes too great the performance may drop even more.

Digital Rapids and Digital Fountain have developed a software simulator to allow you to explore the combined effects of loss and latency You can download the simulator from http://www.digital-rapids.com/Products_Copper.html. Feel free to explore its capabilities.

How competing systems work

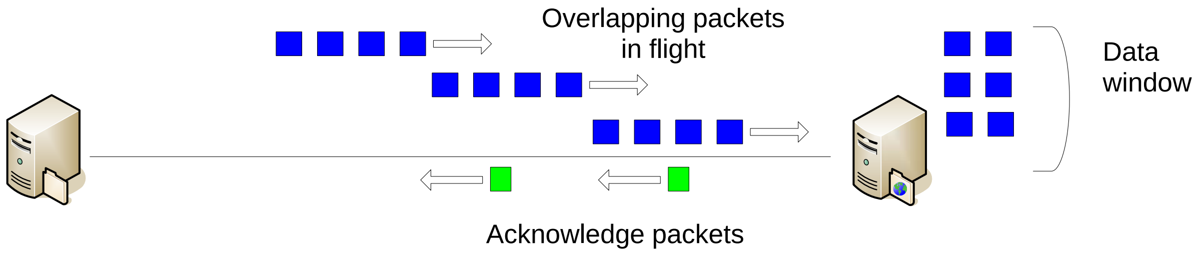

There are TCP/IP based solutions in existence that attempt to overcome loss and latency. One method is to use overlapping transfers either through a TCP/IP proxy or through direct integration with the applications at either end (modified ftp server/client). This use of overlapping transmission can help to overcome latency by having more packets in flight at any given time. The server is not stuck waiting for an acknowledge from the receiver.

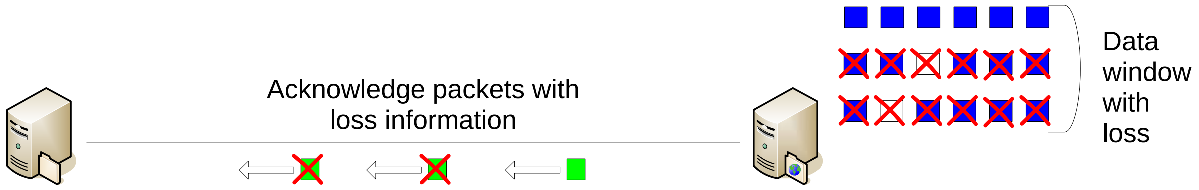

The problem, however, is that loss becomes even more critical as the overlapping transmissions become part of the congestion problem. This can lead to a situation where the bandwidth becomes very difficult to manage the throughput suffers drastically.

The Copper Solution

Copper utilizes Digital Fountain meta-content technology to overcome loss and latency by encrypting the content in such a way that no single packet of data has any particular bearing on the reliability of the stream.

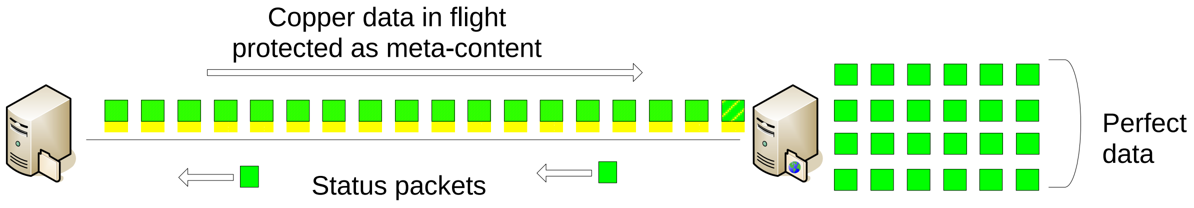

Copper is based on a UDP transport mechanism much like streaming video. Packets are sent without waiting for an acknowledgement from the receiver. This allows the transmitter or sender to utilize a greater percentage of the available network bandwidth.

Unlike TCP/IP based applications, meta-content encoded packets do not need to arrive in any particular order. The data can be reassembled no matter how scrambled the order of the packets on reception.

Copper transmitted packets (meta-content) are also interlinked, and this redundancy allows packets to be lost along the way without any affect on the outcome of the data transmission. The data will arrive 100% intact even if some packets are lost.

A useful analogy for the technology behind Copper is to remember this:

If you had to send a loaf of bread between two points, TCP/IP would send each slice of bread, one piece at a time, in order, bit by bit. If any crumb of bread is missing then the entire slice is not correct and must be resent.

Meta-content instead looks at the loaf of bread, and develops a recipe, not literally of course but figuratively it sends the recipe for creating individual slices of bread and how the loaf fits together. Because of the redundancy built into the algorithm, any loss of packets (recipe for crumbs) along the way can be recreated from another group of packets that have arrived or will arrive during the transmission. In short, Copper sends a much different type of data that is not dependent on packet order or content to be able to recreate with 100% accuracy any type of data it is given.