BIC-TCP

Giving Researchers High-Speed Hopes

The history of TCP (Transmission Control Protocol) shows it's an extremely useful protocol for transmitting data, even over lines with heavy traffic. But long-distance transmissions cause trouble for TCP in the form of long delays over high-speed networks.

With several high-speed, large-bandwidth networks now employed around the world, researchers have worked in recent years to find a new protocol for making full use of the available bandwidth. For example, networks such as ESNet (Energy Sciences Network) offer bandwidth up to 10Gbps when connecting research facilities around the world.

Because these high-speed networks often connect remote locations worldwide, the problems with RTT (round trip time; the time required for a packet to make a trip along with the return response) delays are magnified by the distances the data must travel. Although researchers have introduced and proposed several new protocols in recent years (see the "High-Speed Protocols Compared" sidebar), one of the most promising appears to be BIC-TCP (Binary Increase Congestion Control-TCP). BIC-TCP received the highest marks in a recent series of tests from Stanford researchers on high-speed protocols.

BIC-TCP's Proposal

The trouble with TCP is that it has difficulty filling the bandwidth that's available in a high-speed network (a problem that's usually called congestion control). TCP does increase its congestion window steadily each time it has no packet loss, but the increases are minimal enough that TCP has trouble reaching the full available bandwidth in a desirable time. If it begins encountering errors and packet

Binary Search Increase

The researchers who developed BIC-TCP (Binary Increase Control-Transmission Control Protocol) incorporated binary search increase in the protocol. Binary search increase provides reliable feedback on any network congestion and lost packets, allowing BIC-TCP to aggressively increase its transmission speed toward the maximum allowed by the high-speed network. Binary search increase defines the **minimum window size** as the speed at which no packet loss occurs. The **maximum window size** is the top transmission speed allowed. The **target window size** is the midpoint between the minimum and maximum. (The **current window size** is the actual transmission speed and attempts to climb to the target window size.) Binary search increase calculations cause all of the window sizes to continually change.

1 As the transmission begins, binary search increase calculations allow the current window size to aggressively increase toward the target window size because the difference between the minimum and maximum is large.

2 Once the current window size reaches the target without any packet loss, the minimum window size increases to match the previous target and the target window size increases to the midpoint between the new minimum and maximum (which remains steady). The current window size then continues increasing toward the new target.

3 When a packet loss occurs, the maximum window size drops to match the position of the current window size where the packet loss occurred. The minimum window size drops to match the location where the current window size drops after the packet loss. The target window size is readjusted to the midpoint between the new minimum and maximum. Then the current window size restarts its climb.

4 As the current window size continues to increase without packet losses, binary search increase continues its calculations for the minimum and target window sizes. As the difference between the minimum and maximum window sizes shrinks, binary search increase forces the rate of increase in the current window size to become less aggressive, which allows for fewer packet losses, despite the increase in transmission speed.

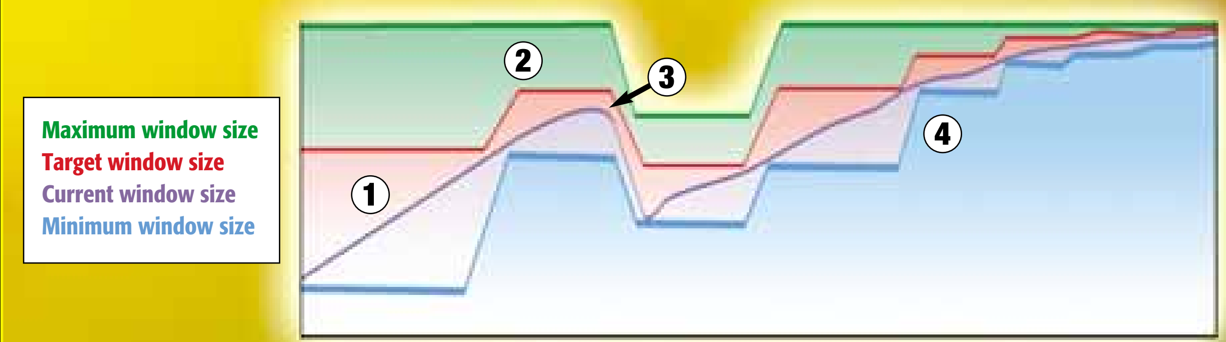

Graphics & Design: Carrie Benes & Jason Codr

TCP's Weakness

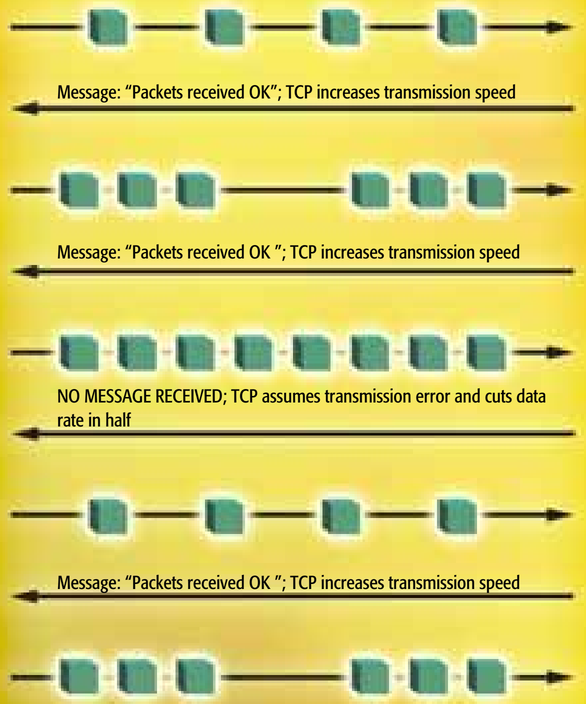

TCP is a key component in transmitting data across computer networks, especially the Internet. Some estimates say TCP accounts for 90% of all data transmission across networks. TCP is extremely reliable because the sending medium always receives notification of the successful transmission of the packets. The reliability factor works especially well on the Internet, where congestion and lost packets are common.

When TCP receives successful notification of the receipt of packets, it slowly increases the transmission speed. However, if TCP doesn't receive a receipt message, it assumes lost packets and network congestion and automatically halves its transmission speed. It then restarts the process of slowly increasing transmission speeds whenever receiving notification of a successful packet receipt.

Most Internet users never notice when TCP slows down to deal with congestion problems because the protocol still works much faster than the typical Internet user's connection speed. TCP will be a reliable, successful protocol for most Internet users for the foreseeable future.

However, in a high-speed network (usually more than 1Gbps in bandwidth), TCP's delays are debilitating. TCP simply cannot take full advantage of the high-speed network's available bandwidth. TCP increases its transmission speeds far too slowly to ever take full advantage of the available bandwidth. Researchers have developed several different Advanced TCP options (such as BIC-TCP) that are more aggressive about increasing transmission speeds while remaining able to work alongside TCP.

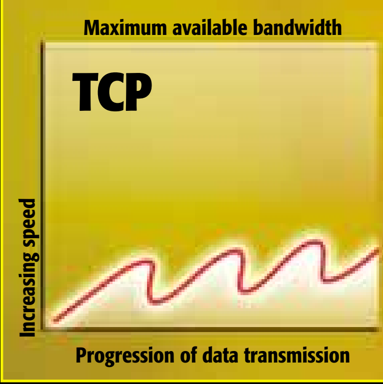

Because of its emphasis on providing reliability, TCP cannot meet the high-speed transmission needs of a large network. As shown at left, TCP's transmission speeds cannot climb quickly enough to meet the maximum because the protocol drops the speed by half each time it encounters an error. Advanced TCP options (shown at right) are more aggressive about increasing transmission speeds, backing off when encountering errors but quickly resuming an aggressive increase in transmission speed.

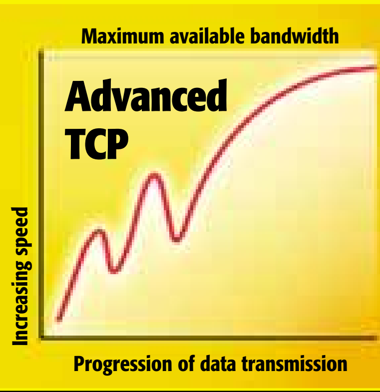

losses, TCP's bandwidth problems become magnified.

Researchers can fix such bandwidth problems with newer protocols designed specifically for high-speed networks. However, one of the biggest concerns researchers have had about some of the high-speed network protocols involves how they work alongside TCP (usually called TCP friendliness). As it works with TCP, the new protocol must treat TCP fairly in the amount of bandwidth it occupies. In other words, while scaling to utilize the full bandwidth of the high-speed network, the new protocol must avoid overwhelming TCP's share of the bandwidth.

The researchers who developed BIC-TCP at North Carolina State University's Department Of Computer Science (Injong Rhee, Lisong Xu, and Khaled Harfoush) had another serious concern about the other available protocols: RTT fairness. When different data flows with varying RTTs occur within a high-speed protocol, some data flows consume more than their share of bandwidth, causing problems. The North Carolina State researchers are hoping they've addressed all of these concerns with BIC-TCP.

When creating BIC-TCP, the researchers put together two important components: Binary search increase and additive increase.

Binary Search Increase. Binary search increase affects the method by which a protocol works to achieve its target sending rate, also called the window. (See the "Binary Search Increase" graphic.) As the protocol works to find the window size that maximizes its bandwidth, the protocol constantly adjusts the current, target, maximum, and minimum window sizes to meet the feedback it receives concerning packet loss.

With TCP, this move toward the maximum available bandwidth occurs far too slowly to achieve the desired results on a high-speed network. When using BIC-TCP, binary search increase allows the protocol to be aggressive when moving from the current window size to the target window size whenever the difference between the two is large. As the difference between the two window sizes shrinks, binary search increase forces the protocol to be less aggressive, which has the effect of treating TCP fairly.

Additive Increase. Additive increase is a process by which TCP sends additional packets per RTT, as long as the previous set of packets arrived without error. BIC-TCP's use of additive increase makes it more stable.

There's No Comparison

One confusing aspect of BIC-TCP has been the widely reported comparisons to DSL that accompanied the recent announcement of the publication of a BIC-TCP research paper. However, because DSL and BIC-TCP are two different items—DSL is a medium and BIC-TCP is a protocol—such comparisons are inaccurate. A protocol actually carries the data across the medium.

Although it's impossible to compare the speed of BIC-TCP with DSL, researchers say BIC-TCP will allow a 10Gbps network to use its full capacity, something TCP alone cannot do. Researchers say that BIC-TCP will provide more benefit to any high-speed network beyond 1Gbps than TCP can. Researchers estimate that over a 10Gbps, long-distance network, TCP only will be able to reach speeds of 450Mbps. When running alongside TCP, BIC-TCP would occupy the remaining bandwidth of the network.

Another source of confusion concerning BIC-TCP concerns how much it will help the average user right now. The answer is not much. Because the vast majority of users connect to the Internet through a cable connection, a DSL connection, or even a dial-up connection, TCP can easily deliver data at the maximum possible bandwidth. TCP's weaknesses really don't show up in a significant way until the protocol is applied to a high-speed network, usually at least 1Gbps. Therefore, most ISPs won't be employing BIC-TCP (or any other high-speed protocol, for that matter) for use with current consumer Internet connections.

Scientists and researchers, however, are ready for the high-speed protocols, which will lead to transmission of petabytes (1 million GB) and exabytes of data (1 billion GB). High-speed protocols will open avenues of research previously thought impossible because of data-transfer limitations. CPU

by Kyle Schurman

High-Speed Protocols Compared

In addition to BIC-TCP (Binary Increase Control-Transmission Control Protocol), several researchers have developed other high-speed protocols–often called Advanced TCP–in recent years. All of these protocols improve on TCP by increasing their congestion windows more quickly than TCP can increase them. These protocols base the rate by which they increase the congestion window size on the overall size of the congestion window, which allows the protocols to maximize the data rate far more quickly than TCP.

Fast TCP. Developed at the California Institute Of Technology, Fast TCP measures a different factor (queuing delay instead of loss probability) to determine congestion in the network. By measuring this factor, Fast TCP can increase its congestion window more quickly than can TCP. Fast TCP was involved in a key test of Advanced TCP over a transatlantic network in 2002.

Hamilton TCP (H-TCP). Developed at the Hamilton Institute in Ireland, H-TCP works at balancing the increased throughput speeds of the high-speed network while maintaining the stability of TCP. H-TCP is aggressive about increasing the congestion window after it reaches the target threshold, therefore more quickly responding to changes in available bandwidth.

High Speed TCP (HS-TCP). HS-TCP uses a table to determine by what factor to increase the congestion window. HS-TCP is similar to H-TCP in its aggressiveness about increasing the congestion window when the available bandwidth increases. However, H-TCP uses an algorithm to determine the rate of increase instead of a table.

High Speed TCP Low Priority (HSTCP-LP). Developed at Rice University, HSTCP-LP only focuses on using the excess network bandwidth instead of attempting to determine its fair share of the bandwidth. Developers of HSTCP-LP borrowed some aspects of HS-TCP.

Scalable TCP (S-TCP). Developed at the Communication Engineering Lab of Cambridge University in England, S-TCP slightly tweaks TCP, allowing it to increase its congestion window exponentially rather than using a linear increase. Because the changes to TCP are minimal, S-TCP researchers say their protocol works extremely well alongside TCP.

The SLAC (Stanford Linear Accelerator Center) recently published its findings on tests it performed on many Advanced TCP options. The SLAC tested each protocol for reliability, fairness, and performance over long-distance networks. SLAC researchers called BIC-TCP the best Advanced TCP performer in its tests, scoring well across all categories. Other protocols, including S-TCP and Fast TCP, also received high marks but had a few minor problems. Researchers say S-TCP was too aggressive in taking bandwidth from TCP and Fast TCP experienced some problems with slowdowns during periods of high reverse traffic. ▲