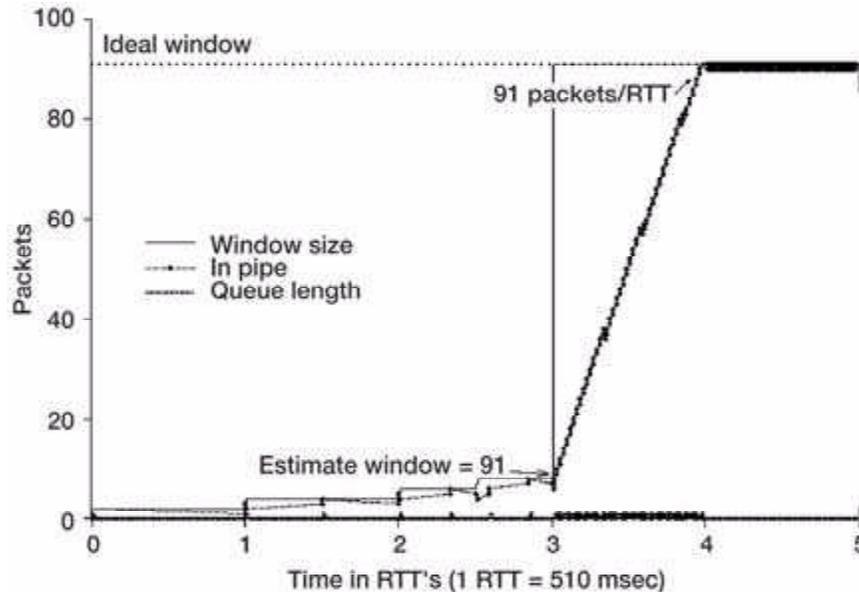


# TCP Pacing Developed

Transmission control protocol (TCP) was conceived and designed to run over a variety of communication links, including wireless and high-bandwidth links. However, with recent technological advances in satellite and fiber-optic networks, researchers are reevaluating the flexibility of TCP. The TCP pacing and packet pair probing implementation may help overcome two of the major obstacles identified for efficient bandwidth utilization over communication links with large delay-bandwidth products.



*Berkeley Network Simulator time-sequence plot showing the possible rapid ascent to full bandwidth utilization when a combination of TCP pacing and packet pair probing is deployed. Single stream over a 500-msec delay link; maximum queue length, 5.*

One problem common to both satellite and fiber-optic networks is that the capacity of these networks, determined by the product of the bandwidth and the delay of the network, can be more than 10 times greater than in conventional networks. The mismatch between the high capacity of these networks and available storage at the intermediate routers in the network poses unique problems for TCP. In a typical network, TCP optimizes its send rate by releasing increasingly large bursts (or windows) of packets (one burst per roundtrip time) to the receiver until it reaches its maximum window size, at which point it has reached the full capacity of the network. In a network with a high delay-bandwidth product, however, TCP's maximum window size may be larger than the queue capacity of some of the network's intermediate routers. Larger windows over-load such router queues, and the routers begin to drop packets. TCP interprets dropped packets as congestion at the bottleneck and reduces its transmission rate, even though no congestion is present. This can result in the ineffective use of available network bandwidth, negatively impacting network performance.

Paced TCP is designed to release its packets smoothly into the network over time rather

than bursting its packets into the network at the transmitting host's network interface line rate. Pacing eliminates the need for excessively large queues throughout the network.

A second problem in obtaining high bandwidth efficiency over large delay-bandwidth links has to do with the slow-start algorithm of TCP. TCP determines the available bandwidth through the use of the control loop, packet send/receipt handshake. As the delay increases, it takes more time for an acknowledgment of each packet. For communications links with large delay-bandwidth products, it can take a relatively long time to reach full bandwidth utilization.

Packet-pair probing is a modification to TCP that attempts to improve network performance by extracting available bandwidth information from the spacing of the acknowledgment packets received at the transmitting host corresponding to the initial packets transmitted.

A modified version of TCP known as TCP pacing with packet pair probing has been developed in FreeBSD and is publicly available for researchers to experiment with and build on. This work was performed by BBN Technologies through the Professional, Administrative, Computing and Engineering Contract at the NASA Glenn Research Center. The software is available online at <http://www.ir.bbn.com/projects/pace/>, with the "readme" file available separately.

## References

1. Partridge, Craig; Shepard, Tim; and Coulter, Robert: Study and Simulation of Enhancements for TCP Performance Over Noisy High Latency Links. NASA/CR-1999-209167, 1999.
2. Partridge, Craig; Sterbenz, James; and Rockwell, Dennis: TCP Pacing/Packet Pair Probing. BBN Technologies, Cambridge, MA. Also NASA TM (to be published in 2002).

**Glenn contact:** William D. Ivancic, 216-433-3494, William.D.Ivancic@grc.nasa.gov

**Author:** William D. Ivancic

**Headquarters program office:** OSF (SOMO)

**Programs/Projects:** TSIS