CONSTELLATION
CREATING RELATIONSHIPS

Optimized TCP/IP for Satellite Communication

Satellites seem an ideal means for offering Internet and intranet access over long distances and to remote locations. However, Internet protocols are not optimized for satellite conditions, and consequently, the throughput over satellite networks is restricted to only a fraction of the available bandwidth. iDirect engineering has addressed and overcome these issues. The TCP/IP protocol was not designed to perform well over high-latency or noisy channels. Geostationary satellite links are inherently slow and can be noisy. TCP/IP based data communication is almost useless at Bit Error Rates (BER) of 10^{-7} or higher. The TCP/IP shortcomings in the typical satellite environment—degradations due to slow-start, window size, and acknowledgment frequency—are well known. There have been attempts to deliver IP over satellite, but the satellite technologies have focused on connection-oriented transmission protocols that are really suited for voice traffic rather than IP, and unnecessarily waste expensive capacity.

The iDirect development team has designed an innovative technique to overcome these IP-over-satellite technology difficulties. The iDirect solutions are presented in the following sections.

Performance of TCP/IP over Satellite – Latency Issues

Communications over geo-synchronous satellites have round trip times of approximately 560 ms. The journey through the atmosphere can also introduce bit errors into the data stream. These factors, combined with back channel bandwidth typically much smaller than that available on the forward channel, reduce the effectiveness of TCP, which is optimized for short hops over low-loss cable or fiber. Satellite conditions adversely affect a number of elements of the TCP architecture, including its congestion avoidance algorithms, data acknowledgment mechanisms, and window size limitations, which combine to severely constrict the data throughput rate that can be achieved over satellite links.

Congestion Avoidance: In order to avoid the possibility of congestive network meltdown, TCP assumes that all data loss is caused by congestion and responds by reducing the transmission rate using a feature called “slow-start”, to discover the throughput capacity upon initial connection setup. Slow-start sends a packet across the physical connection and waits for a response. If a response is received, the next packet is sent a bit faster. This procedure is repeated until the speed of the link is discovered. With the half-second delay between responses, throughput is significantly slowed. Clearly, if slow-start can be bypassed, a significant drag on TCP/IP performance can be removed.

Data Acknowledgements: Latency is exacerbated in TCP/IP transmissions by the fact that the protocol requires acknowledgements for packets sent across the link. The simple, heuristic data acknowledgment scheme used by TCP does not adapt well to long latency or highly asymmetric bandwidth conditions. To provide reliable data transmission, the TCP receiver constantly sends acknowledgments for the data received back to the sender. This is to ensure reliable communication under uncertain and congested network conditions. However, the acknowledgment mechanism becomes a problem for short transmissions over high latency channels. Each packet exchange takes at least ~560 mSecs to complete. While this may be a smaller factor in broadcast-only applications, it can significantly limit throughput for typical interactive IP applications such as Internet, Intranet, and Extranet.

Window Size: TCP/IP has a built-in windowing mechanism that determines the highest possible throughput, while balancing the risk of re-transmission of dropped packets. It works by allowing a transmitter to send a number of packets before having to wait for an acknowledgment from the receiver. Typically, the window size is set to accommodate low-latency (terrestrial) connections. So the windows tend to be short with respect a satellite link, but minimize the amount of data that would need to be retransmitted when packets are dropped. Short windows slow down a satellite connection dramatically. Opening up the constricted flow of packets requires adjusting (tuning) window size for the known latency and expected noise performance of the link.

iDirect's TCP Acceleration Protocol and Lower Bit Error Rates Over Latency Issues

To overcome slow-start, iDirect has developed TCP Acceleration protocol to work over the satellite link. The iDirect system has reduced Bit Error Rates to rates compatible with good TCP/IP performance.

In the iDirect network all TCP connections are terminated and re-originated at both ends of the satellite link. Over the satellite link, a new protocol is used that is both transparent to TCP/IP and optimized for the satellite environment. This satellite link protocol does not use slow-start. Since iDirect controls the link and schedules capacity, the throughput for a TCP connection is known prior to transmission and no discovery process is required. The data is sent at a predetermined rate over the link without any round-trip delays.

The characteristics of BER and latency on the satellite link are known, so an appropriate window size can be selected. Even under heavy rain fade conditions, the iDirect system will perform at ~10^{-9} BER. Under clear sky conditions, the iDirect system will perform at 10^{-10} BER or better. This combined with the systems capability to automatically adjust power budget in a link, keeps the BER at ~10^{-9} or better. At these performance levels, very large window sizes are possible, which effectively removes throughput degradation due to window size. The low BER also allows a reduced frequency of acknowledgments, for increased network efficiency, without sacrificing reliability. The low Bit Error Rate of the iDirect network and resulting reduced acknowledgment frequency contribute to reducing the drag on TCP/IP performance due to latency. The iDirect system operates at ~80% of link capacity, this is due to the advanced TPC (Turbo Product Code) ECC mechanism that is implemented.

iDirect Protocol Design

At the heart of the iDirect system is the iDirect Protocol, optimized to provide maximum throughput for satellite networks. The iDirect Protocol is designed to respond efficiently to typical satellite latency, bit errors, and asymmetric bandwidth conditions, and to take advantage of optimizations possible on a single-path link with known bandwidth.

Efficient Acknowledgment Algorithm: The iDirect Protocol utilizes a highly efficient selective retransmission algorithm for the acknowledgment of data. Because there is only a single-path over the satellite for all packets with no intermediate routing, any gaps in the packet sequence can be assumed to be data loss due to corruption rather than network congestion. The receiving NetModem immediately requests retransmission of the missing data from the transmitting HubModem.

Because the iDirect Protocol does not use acknowledgments as the primary means of identifying lost data, it needs only infrequent acknowledgments to confirm data arrival and clear buffers. In contrast, TCP sends a constant stream of acknowledgments over the reverse channel. The iDirect Protocol reduces back channel usage by 70% or more, thereby dramatically increasing the performance of networks where limited back channel bandwidth is the system bottleneck.

Large Windows: The iDirect Protocol offers essentially unlimited window sizes for transmission of data between the Modems. Because the bandwidth-delay product over the satellite between the iDirect Modems is much larger than that from the NetModem to the end node, the large iDirect Protocol window effectively removes the dependency of the network on the bandwidth-delay product.

Congestion Avoidance: The iDirect Protocol does not use unnecessary congestion avoidance mechanisms for the hop over the satellite between the iDirect Modems. However, the system continues to use Slow Start and all other standard TCP congestion avoidance algorithms for the terrestrial connections between the iDirect Modems and the end nodes. The iDirect Protocol also uses a streamlined handshake mechanism to reduce the number of round-trip times required for connection set-up.

TCP Acceleration Protocol – How it works

When a session is initiated (SYN) from the remote side and a response is received from the receiving host, the HubModem (at the hub site), sends a ACK back to the host, to keep sending packets to it. In the meanwhile, the HubModem sends the packet to the remote. The HubModem in essence handles the TCP session with the host and a session with the remote. This method enables the capability of handling lost packets on either side, which most other implementations do not provide. The figure below shows the basic operation of TCP Acceleration.