



Report ITU-R S.2148
(09/2009)

Transmission control protocol (TCP) over satellite networks

S Series
Fixed-satellite service



Documents
Document Sharing Hub

Foreword

The role of the Radiocommunication Sector is to ensure the rational, equitable, efficient and economical use of the radio-frequency spectrum by all radiocommunication services, including satellite services, and carry out studies without limit of frequency range on the basis of which Recommendations are adopted.

The regulatory and policy functions of the Radiocommunication Sector are performed by World and Regional Radiocommunication Conferences and Radiocommunication Assemblies supported by Study Groups.

Policy on Intellectual Property Right (IPR)

ITU-R policy on IPR is described in the Common Patent Policy for ITU-T/ITU-R/ISO/IEC referenced in Annex 1 of Resolution ITU-R 1. Forms to be used for the submission of patent statements and licensing declarations by patent holders are available from <http://www.itu.int/ITU-R/go/patents/en> where the Guidelines for Implementation of the Common Patent Policy for ITU-T/ITU-R/ISO/IEC and the ITU-R patent information database can also be found.

Series of ITU-R Reports

(Also available online at <http://www.itu.int/publ/R-REP/en>)

Series	Title
BO	Satellite delivery
BR	Recording for production, archival and play-out; film for television
BS	Broadcasting service (sound)
BT	Broadcasting service (television)
F	Fixed service
M	Mobile, radiodetermination, amateur and related satellite services
P	Radiowave propagation
RA	Radio astronomy
RS	Remote sensing systems
S	Fixed-satellite service
SA	Space applications and meteorology
SF	Frequency sharing and coordination between fixed-satellite and fixed service systems
SM	Spectrum management

***Note:** This ITU-R Report was approved in English by the Study Group under the procedure detailed in Resolution ITU-R 1.*

Electronic Publication
Geneva, 2010

© ITU 2010

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without written permission of ITU.

REPORT ITU-R S.2148

Transmission control protocol (TCP) over satellite networks

(2009)

TABLE OF CONTENTS

	<i>Page</i>
1 Introduction	3
2 Satellite system reference models.....	3
2.1 Point-to-point links	3
2.2 VSAT networks	3
2.2.1 Star topology	3
2.2.2 Mesh topology.....	4
2.3 Broadband access.....	4
3 TCP limitations over satellite links	5
3.1 Bandwidth-delay product.....	5
3.2 Slow start and congestion avoidance	6
3.3 Fast retransmit and fast recovery	7
3.4 Effect of bit errors on TCP throughput.....	8
3.5 Asymmetric networks.....	8
3.6 Advice to designers.....	9
4 TCP enhancement methodologies	9
4.1 Variations of baseline TCP	9
4.2 Segment splitting methodologies.....	16
4.2.1 Two-segment splitting methodology	16
4.2.2 Three-segment splitting methodology	17
4.2.3 Discussion	19
4.3 Caching and spoofing	19
4.3.1 Caching	19
4.3.2 Spoofing.....	19
4.3.3 Spoofing and caching.....	19
5 Performance enhancing proxies	20
5.1 TCP spoofing	21
5.2 TCP splitting.....	21
5.3 Other PEP mechanisms.....	21
5.4 Implications of using PEP	22
5.4.1 End-to-end security	22
5.4.2 End-to-end reliability	22
6 Other transmission control protocols.....	22
6.1 Space communication protocol specification – transport protocol (SCPS-TP).....	22
6.2 Xpress transport protocol.....	22
6.3 Stream control transmission protocol	23
6.4 Comparison between transmission control protocols	24
7 Conclusion.....	24

List of acronyms

ACK	Acknowledgement
ATM	Asynchronous transfer mode
ARQ	Automatic repeat request
BDP	Bandwidth-delay product
BER	Bit-error ratio
cwnd	Congestion window (variable in TCP)
DACK	Delayed acknowledgement
ECN	Explicit congestion notification
FIN	Final segment (in a TCP connection)
GSO	Geostationary-satellite orbit
GW	Gateway
HEO	High-Earth orbit
IETF	Internet engineering task force
IP	Internet protocol
IPSEC	IP security protocol
ISP	Internet service provider
LAN	Local area network
LEO	Low-Earth orbit
MEO	Medium-Earth orbit
MPLS	Multiprotocol label switching
MTU	Maximum transmission unit
PAWS	Protect against wrapped sequence(s)
PEP	Performance enhancing proxy
RAM	Random access memory
RBP	Rate-based pacing
RFC	Request for comments (issued by the IETF)
RTT	Round trip time
RTTM	RTT measurement
SACK	Selective acknowledgment
SCPS-TP	Space communication protocol specification – transport protocol
SCTP	Stream control transmission protocol
SYN	Synchronous start segment (used to establish a TCP connection)
T/TCP	TCP for transactions
TCP	Transmission control protocol
UDP	User datagram protocol
VSAT	Very small aperture terminal
XTP	Express transport protocol


1 Introduction

This Report presents reference models of networks including a satellite link, to carry IP packets, followed by a description of the limitations of transmission control protocol (TCP) over satellite links, as well as, various methodologies to overcome the limitations.

2 Satellite system reference models

2.1 Point-to-point links

Figure 1 provides a reference model for a network carrying IP packet transmissions. The network consists of a satellite link and associated terrestrial networks between two end-users. The satellite link is bidirectional and consists of link AB (from earth station A to earth station B with an information bit rate, R_{AB}) and of link BA (from earth station B to earth station A with an information bit rate, R_{BA}). The terrestrial networks can employ various data link layer protocols (e.g., asynchronous transfer mode (ATM), frame relay, MPLS).




NOTE 1 – The reference model above considers only one satellite hop. Throughout this Report the techniques that segment the TCP connection to improve TCP performance over satellite links are described for one satellite hop. However an end-to-end connection may include several satellite hops. In this case, such techniques will have to be implemented over each individual satellite link.

2.2 VSAT networks

2.2.1 Star topology

Figure 2 depicts the standard star configuration in which signals from various remote users connect to a gateway earth station which in turn connects to terrestrial network.

FIGURE 2
Star topology




Report 2148-02

2.2.2 Mesh topology

Figure 3 illustrates a mesh configuration whereby any pair of earth stations can be connected directly via satellite.

FIGURE 3
Mesh topology



Report 2148-03

2.3 Broadband access

Even if not completely similar to very small aperture terminal (VSAT) networks, broadband access networks generally use the same topologies (i.e. star or mesh).

3 TCP limitations over satellite links

The TCP cannot distinguish the performance degradation caused by link errors from congestion. It assumes that any loss in the network is due to congestion only and the sender responds by reducing its packet transfer rate.

The baseline TCP (TCP Reno) specifies slow start, congestion avoidance, fast retransmit and fast recovery for congestion control. The TCP uses window flow control mechanism in which the transmission window allows the receiving TCP to control the amount of data being sent to it at any given time. The receiver advertises a window size to the sender. The window measures, in byte/s, the amount of unacknowledged data that the sender can have in transit to the receiver.

3.1 Bandwidth-delay product

The bandwidth-delay product (BDP) defines the amount of data a TCP connection should have “in flight” (data that has been transmitted, but not yet acknowledged) at any time to fully utilize the available channel capacity. The delay is the round trip time (RTT) and the bandwidth is the capacity of the bottleneck link in the path.

For links with a large BDP, such as in geostationary satellite networks, TCP senders and receivers with limited congestion/receive windows will not be able to take advantage of the available bandwidth. The standard maximum TCP window of 65 535 byte/s is not adequate to allow a single TCP connection to utilize the entire bandwidth available on some satellite channels. In a loss-free network the TCP throughput is limited by equation (1):

$$\text{Maximum throughput} = \frac{\text{Window size}}{\text{RTT}} \quad (1)$$

Therefore, when using the maximum TCP window size of 64 kbyte/s and satellite links with variable RTT, the maximum throughput is as follows:

TABLE 1
Maximum throughput according to RTT values

Satellite network type	RTT (ms)	Maximum throughput (kbyte/s)
LEO	~20	~3 200
MEO	~200	~320
HEO	~600	~110
GSO	~520	~120

NOTE 1 – The above-mentioned RTT do not take into account any buffer delay but are computed on the basis of the propagation delay.

IETF¹ Request for Comments 3150² recommends performance of network paths that traverse “very low bit-rate” links. It is applicable for any network where hosts can saturate available bandwidth, but the design space explicitly includes connections that traverse 56 kbit/s modem links or 4.8 kbit/s

¹ Internet Engineering Task Force.

² End-to-end Performance implications of slow links.

wireless access links. Some of the discussion is common with Request for Comments 2689³, using header compression. It focuses more on traditional data applications for which “best-effort delivery” is appropriate.

3.2 Slow start and congestion avoidance


The TCP sender maintains a congestion window to measure the network capacity. The number of unacknowledged packets in the network is limited to this value (or to the receiver advertised window whichever is lower). At the start of a TCP connection, the congestion window is set to one TCP segment. It increases by one segment on the receipt of each new acknowledgment until it reaches its maximum value of 64 kbyte/s. The sender maintains a retransmission time out for the last unacknowledged packet. Congestion is detected by the expiration of the retransmission time out. When the timer expires, the sender saves the value of half the congestion window (called slow start threshold) and sets it to one segment. The sender then retransmits segments starting from the lost segment. The congestion window is increased by one segment on the receipt of each new acknowledgment until it reaches the slow start threshold. This is the slow start phase. After that, the congestion window increases by one segment every RTT. This results in a linear increase of the congestion window every RTT and is called the congestion avoidance phase. Figure 4 shows the slow start and congestion avoidance phases for a typical TCP connection (in Fig. 4, cwnd stands for congestion window).

The time required by the slow start mechanism to reach a bit rate B is given by equation (2):

$$\text{Slow start duration} = \text{RTT} \left(1 + \log_2 \frac{B \cdot \text{RTT}}{l} \right) \quad (2)$$

where l is the average packet length expressed in bits.

Table 2 shows the duration of slow start phase for various satellite orbits and different values of bit rates B , when $l = 1$ kbit.



³ RFC 2689 “Providing integrated services over low-bit rate links”.