

Internet Engineering Task Force
INTERNET DRAFT
File: [draft-allman-tcp-sack-00.txt](#)

Mark Allman
NASA GRC/BBN
Ethan Blanton
Ohio University
November, 2000
Expires: May, 2001

A Conservative SACK-based Loss Recovery Algorithm for TCP

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This document presents a conservative loss recovery algorithm for TCP that is based on the use of the selective acknowledgment TCP option. The algorithm presented in this document conforms to the spirit of the current congestion control specification, but allows TCP senders to recover more effectively when multiple segments are lost from a single flight of data.

1 Introduction

This document presents a conservative loss recovery algorithm for TCP that is based on the use of the selective acknowledgment TCP option. While the TCP selective acknowledgment (SACK) option [[RFC2018](#)] is being steadily deployed in the Internet [[A1100](#)] there is evidence that hosts are not using the SACK information when making retransmission and congestion control decisions [[PF00](#)]. The goal of this document is to outline one straightforward method for TCP implementations to use SACK information to increase performance.

[[RFC2581](#)] allows advanced loss recovery algorithms to be used by TCP

[[RFC793](#)] provided that they follow the spirit of TCP's congestion control algorithms [RFC2581, [RFC2914](#)]. [[RFC2582](#)] outlines one such advanced recovery algorithm called NewReno. This document outlines

Expires: May 2001

[Page 1]

a loss recovery algorithm that uses the selective acknowledgment (SACK) [[RFC2018](#)] TCP option to enhance TCP's loss recovery. The algorithm outlined in this document, heavily based on the algorithm detailed in [[FF96](#)], is a conservative replacement of the fast recovery algorithm [[Jac90](#), [RFC2581](#)]. The algorithm specified in this document is a straightforward SACK-based loss recovery strategy that follows the guidelines set in [[RFC2581](#)] and can safely be used in TCP implementations. Alternate SACK-based loss recovery methods can be used in TCP as implementers see fit (as long as the alternate algorithms follow the guidelines provided in [[RFC2581](#)]).

2 Definitions

The reader is expected to be familiar with the definitions given in [[RFC2581](#)].

For the purposes of explaining the SACK-based loss recovery algorithm we define two variables that a TCP sender stores:

``HighACK'' is the sequence number of the highest cumulative ACK received at a given point.

``HighData'' is the highest sequence number transmitted at a given point.

For the purposes of this specification we define a ``duplicate acknowledgment'' as an acknowledgment (ACK) whose cumulative ACK number is equal to the current value of HighACK and also conveys new selective acknowledgment information for segment(s) above HighACK.

3 Keeping Track of SACK Information

For a TCP sender to implement the algorithm defined in the next section it must keep a data structure to store incoming selective acknowledgment information on a per connection basis. Such a data structure is commonly called the ``scoreboard''. For the purposes of the algorithm defined in this document the scoreboard MUST implement the following functions:

Update ():

Each octet that is cumulatively ACKed or SACKed should be marked accordingly in the scoreboard data structure, and the total number of octets SACKed should be recorded. Note that SACK information is advisory and therefore SACKed data MUST NOT be removed from TCP's retransmission buffer until the data is cumulatively acknowledged.

MarkRetran ():

When a retransmission is sent, the scoreboard MUST be updated with this information so that data is not repeatedly retransmitted by the SACK-based algorithm outlined in this document. Note: If a retransmission is lost it will be repaired

Expires: May 2001

[Page 2]

using TCP's retransmission timer.

NextSeg ():

This routine MUST return the sequence number range of the oldest segment that has not been cumulatively ACKed or SACKed and not been retransmitted. If no such segment is available this routine MUST return the sequence number range for the first previously unsent segment (if such a segment exists).

AmountSACKed ():

This routine MUST return the number of octets selectively acknowledged by the receiver.

LeftNetwork ():

This function MUST return the number of octets in the given sequence number range that have left the network. The algorithm checks each octet in the given range and separately keeps track of the number of retransmitted octets and the number of octets that are cumulatively ACKed but were not SACKed. Note: it is possible to have octets that fit both categories. In this case, the octets MUST be counted in both categories. After checking the sequence number range given this routine returns the sum of the two counters.

Note: The SACK-based loss recovery algorithm outlined in this document requires more computational resources than previous TCP loss recovery strategies. However, we believe the scoreboard data structure can be implemented in a reasonably efficient manner (both in terms of computation complexity and memory usage) in most TCP implementations.

4 Algorithm Details

Upon the receipt of the first and second duplicate ACKs, the scoreboard MUST be updated per the selective acknowledgment information contained in the ACK (via the Update () routine). Note: The first and second duplicate ACKs can also be used to trigger the transmission of previously unsent segments using the Limited Transmit mechanism [[ABF00](#)].

When a TCP sender receives the third duplicate ACK the scoreboard MUST be updated with the new SACK information (via Update ()) and a loss recovery phase SHOULD be initiated, per the fast retransmit algorithm outlined in [[RFC2581](#)], and the following steps MUST be taken:

(1) Set a ``pipe'' variable to the number of outstanding octets

(i.e., octets that have been sent but not yet acknowledged), per the following equation:

$$\text{pipe} = \text{HighData} - \text{HighACK} - \text{AmountSACKed} ()$$

Expires: May 2001

[Page 3]

- (2) Set a ``RecoveryPoint'' variable to HighData. When the TCP sender receives a cumulative ACK for this data octet the loss recovery phase is terminated.
- (3) The congestion window (cwnd) is reduced to half its current value. The value of the slow start threshold (ssthresh) is set to the halved value of cwnd.
- (4) Retransmit the first data segment not covered by HighACK. Use the MarkRetran () function to mark the sequence number range as having been retransmitted in the scoreboard.

Once a TCP is in the loss recovery phase the following procedure MUST be used for each arriving ACK:

- (A) An incoming cumulative ACK for a sequence number greater than or equal to RecoveryPoint signals the end of loss recovery and the loss recovery phase MUST be terminated.
- (B) Upon receipt of a duplicate ACK the following actions MUST be taken:
 - (B.1) Use Update () to record the new SACK information conveyed by the incoming ACK.
 - (B.2) The pipe variable is decremented by the number of newly SACKed data octets conveyed in the incoming ACK, as that is the amount of new data that has left the network.
- (C) When a ``partial ACK'' (an ACK that increases the HighACK point, but does not terminate loss recovery) arrives, the following actions MUST be performed:
 - (C.1) Before updating HighACK based on the received cumulative ACK, save HighACK as OldHighACK.
 - (C.2) The scoreboard MUST be updated based on the cumulative ACK and any new SACK information that is included in the ACK via the Update () routine.
 - (C.3) The value of pipe MUST be decremented by the number of octets returned by the LeftNetwork () routine when given the sequence number range OldHighACK-HighACK.
- (D) If pipe is less than cwnd and the receiver's advertised window permits, the TCP sender SHOULD transmit a segment, as follows:
 - (D.1) The scoreboard MUST be queried via NextSeg () for the sequence number range of the next segment to transmit, and

the given segment is sent.

(D.2) The pipe variable MUST be incremented by the number of data octets sent in (D.1).

Expires: May 2001

[Page 4]

5 Research

The algorithm specified in this document is analyzed in [FF96], which shows that the above algorithm is effective in reducing transfer time over standard TCP Reno [RFC2581] when multiple segments are dropped from a window of data (especially as the number of drops increases). [AHK097] shows that the algorithm defined in this document can greatly improve throughput in connections traversing satellite channels.

6 Security Considerations

The algorithm presented in this paper shares security considerations with [RFC2581]. A key difference is that an algorithm based on SACKs is more robust against attackers forging duplicate ACKs to force the TCP sender to reduce cwnd. With SACKs TCP senders have an additional check on whether the ACK is legitimate or not. While not fool-proof, SACK provides some amount of protection in this area.

Acknowledgments

The authors wish to thank Sally Floyd for encouraging this document and commenting on an early draft. The algorithm described in this document is largely based on an algorithm outlined by Kevin Fall and Sally Floyd in [FF96] (although the authors of this document assume responsibility for any mistakes in the above). We thank Vern Paxson for providing valuable feedback on an early version of this draft. Finally, we thank Matt Mathis and Jamshid Mahdavi for implementing the scoreboard in ns and hence guiding our thinking in keeping track of SACK state.

References

- [ABF00] Mark Allman, Hari Balakrishnan, Sally Floyd. Enhancing TCP's Loss Recovery Using Limited Transmit, August 2000. Internet-Draft [draft-ietf-tsvwg-limited-xmit-00.txt](#) (work in progress).
- [AHK097] Mark Allman, Chris Hayes, Hans Kruse, Shawn Ostermann. TCP Performance Over Satellite Links. Proceedings of the Fifth International Conference on Telecommunications Systems, Nashville, TN, March, 1997.
- [All00] Mark Allman. A Web Server's View of the Transport Layer. ACM Computer Communication Review, 30(5), October 2000.
- [FF96] Kevin Fall and Sally Floyd. Simulation-based Comparisons of Tahoe, Reno and SACK TCP. Computer Communication Review, July 1996.

[Jac90] Van Jacobson. Modified TCP Congestion Avoidance Algorithm.
Technical Report, LBL, April 1990.

Expires: May 2001

[Page 5]

- [PF00] Jitendra Padhye, Sally Floyd. TBIT, the TCP Behavior Inference Tool, October 2000. <http://www.aciri.org/tbit/>.
- [RFC793] Jon Postel, Transmission Control Protocol, STD 7, [RFC 793](#), September 1981.
- [RFC2581] Mark Allman, Vern Paxson, W. Richard Stevens, TCP Congestion Control, [RFC 2581](#), April 1999.
- [RFC2582] Sally Floyd and Tom Henderson. The NewReno Modification to TCP's Fast Recovery Algorithm, [RFC 2582](#), April 1999.
- [RFC2914] Sally Floyd. Congestion Control Principles, [RFC 2914](#), September 2000.

Author's Addresses:

Mark Allman
NASA Glenn Research Center/BBN Technologies
Lewis Field
21000 Brookpark Rd. MS 54-2
Cleveland, OH 44135
Phone: 216-433-6586
Fax: 216-433-8705
mallman@grc.nasa.gov
<http://roland.grc.nasa.gov/~mallman>

Ethan Blanton
Ohio University Internetworking Research Lab
Stocker Center
Athens, OH 45701
eblanton@cs.ohiou.edu

Expires: May 2001

[Page 6]