

Eliminating Aborted Data Delivery Over Cellular Links

Andrei Gurtov
University of Helsinki/ICSI
gurtov@icsi.berkeley.edu

EXTENDED ABSTRACT

Internet users often abort their transfers in progress, for example by clicking on the "Reload" button or another link in a web browser. Analysis of backbone Internet traces [5] in Table 1 shows that 15-30 % of all TCP connections are abnormally terminated via a reset. The average length of reset connections is not significantly different from completed connections. Thus, resets are not merely generated by TCP connection refusals [1].

Packets from aborted transport connections are often sent unnecessarily to the user over a slow last-hop link delaying useful traffic. This is a particular concern for wide-area wireless links, because unnecessary transmissions waste scarce radio bandwidth, battery power at the mobile terminal and incur monetary cost due to charging by data volume.

Although the total fraction on aborted traffic is only a few percent in backbone Internet traces in Table 1, we believe it is a significant problem for slow last-hop links for two reasons. First, users easily get impatient while waiting for a transfer to complete over a slow link. Second, the round trip time for such links can reach several seconds during a bulk transfer. This delays delivery of an abort notification from the receiver to the sender.

Measurements in a live cellular network show that loading of a new web page is often delayed up to ten seconds due to delivery of data from previously aborted pages. Figure 1 shows a trace of a web page download being aborted. It takes seven seconds until the server receives the abort notification, stops sending and page data drains from the link queue. About 25 kilobytes of data is unnecessarily transmitted over a 30-kbps GPRS link.

Deploying an active queue management algorithm, such as RED, in the last-hop router can reduce the number of aborted packets sent over the last-hop link.

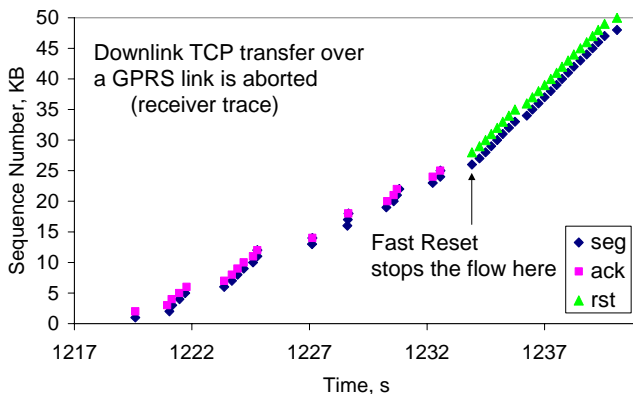


Figure 1. Aborted TCP connection over GPRS to www.zed.com (Mozilla/HTTP 1.1).

AQM keeps the average queue size low without penalizing bursty sources. However, cellular links require a buffer larger than the bandwidth-delay product of the link to efficiently implement local error recovery. In existing cellular networks we often find the bottleneck buffer of about ten times the bandwidth-delay product of the link [3].

The standard TCP receiver generates reset (RST) packets after receiving (and discarding) packets on an aborted connection. We propose a Fast Reset algorithm to eliminate delivery of aborted data over a last-hop link. The algorithm can be included into a performance enhancing proxy snooping packet headers at the last-hop link. The algorithm is illustrated in Figure 2 and has following steps:

1. An application on the mobile client opens a TCP connection and requests a web object.
2. The server receives the request and starts transmitting data to the client.
3. Data packets are buffered in the last-hop router and transmitted to the client.
4. The user decides to abort the download, for example by pressing a 'Reload' button. The TCP receiver responds with RST packets to data packets arriving from the server.
5. The last-hop router notices a RST packet and discards buffered packets in the downlink direction that belong to the aborted connection. It then forwards the RST packet toward the server.
6. The server receives the RST packet and stops transmitting data on the aborted connection.

Step 5 can be further detailed as follows:

```
for each arriving pkt1 in uplink
  if pkt1 is RST
    for each queued pkt2 in downlink
      if (src1, dst1, sport1, dport1, prot1)
        == (dst2, src2, dport2, sport2, prot2)
        discard(pkt2)
forward(pkt1)
```




Figure 2. The Fast Reset algorithm.

Table 1. Resets in the Internet traffic.

Trace id	TcpCons	RstCons	RstCons/ TcpCons	Pkts/ TcpCons	Pkts/ RstCons	TotalPkts	RstPkts	2*RstPkts/ TotalPkts
200011091400	40000	8520	0.21	23.37	14.51	1999982	12528	0.013
200206121400	184000	56379	0.31	17.06	10.16	4901552	71766	0.029
200301051400	143000	17608	0.12	20.36	24.47	4979639	34742	0.014
200201121400	120000	32952	0.27	18.30	11.47	2993417	61208	0.041
200301021400	135000	20606	0.15	18.33	14.17	4976714	39402	0.016

The 5-tuple (src, dst, sport, dport, prot) uniquely identifies a transport connection by its source and destination IP addresses, source and destination port numbers and the protocol number.

If a RST packet gets lost, the sender retransmits a TCP segment after a timeout. This segment will not be dropped by the last-hop router according to the algorithm. Instead, the retransmitted segment will generate another RST packet at the receiver. There is no state kept in the router, thus there is no harm to new TCP connections if the client reuses ports from aborted connections. Fast Reset works robustly for all TCP applications including HTTP, FTP, and peer-to-peer.

Since web browsing is the main contributor of aborted connections, we describe how resets affect HTTP in detail. HTTP v1.0 allows maximum of four concurrent TCP connections to a server. Thus, aborting a web page download can generate resets for several TCP connections. This reduces performance benefit of Fast Reset, because aborted data contained in the buffer of last-hop router is distributed over several TCP connections.

HTTP v1.1 defines persistent connections. With pipelining several HTTP requests can be outstanding over a single TCP connection. Once an abort of a single request is initiated, the whole pipeline is aborted. This helps the Fast Reset algorithm, because with a single reset packet unnecessary data from several data objects can be eliminated.

We evaluated Fast Reset using the Netscape browser in Linux over a GPRS cellular link. When loading of a web page is aborted, 5-25 packets unnecessarily arrive to the receiver from aborted TCP connections. This corresponds to 1-10 seconds of time wasted before a new page begins loading. Fast Reset reduces the penalty of aborting a web page to one-two packets per a TCP connection. As a result, response time can be improved by 20 %, battery power and radio spectrum are preserved.

The main limitation of the Fast Reset algorithm is the requirement for a TCP connection to traverse through the last-hop router in both directions. Although we expect this requirement to be met in most cases, there are asymmetric setups, e.g. for satellite links [7]. Another drawback of Fast Reset is processing load on the router to parse the queue on every arriving RST packet. We expect the load to be acceptable for low rate links [2]. Finally, Fast Reset creates a layering violation because the router has to access transport-layer headers. TCP/IP header compression is another example of a layering violation. If

the network-layer protocol, such as ISO CONP, is connection-oriented, Fast Reset can avoid this problem.

Fast Reset can be also implemented in the mobile client to eliminate aborted data delivery for uplink transfers. This would add little additional benefit because mostly the client and not servers initiates aborts. Fast Reset can be adopted for other connection-oriented transport protocols such as SCTP [6] and DCCP [4].

Generally we feel that aborts of transport connections have not received sufficient attention in networking research. Fingerprinting of packets was proposed to eliminate redundant data delivery [8], but this study is orthogonal to the problem of unnecessary data delivery from aborted connections. We are not aware of any HTTP traffic generator that would take into account aborting of web page transfers. Typically, such generators include 'thinking time' between loading of web pages. If the user starts loading a new page before the old one is completely loaded, then thinking time becomes 'negative'.

REFERENCES

- [1] S. Floyd, Inappropriate TCP Resets Considered Harmful, IETF RFC 3360, August 2002.
- [2] P. Gupta, N. McKeown, Packet classification on multiple fields, In Proc. ACM SIGCOMM'99, August 1999.
- [3] A. Gurtov, M. Passoja, O. Aalto, M. Raitola, Multi-Layer Protocol Tracing in a GPRS Network, In Proc. IEEE VTC'02, September 2002.
- [4] E. Kohler, M. Handley, S. Floyd, J. Padhye, Datagram Congestion Control Protocol (DCCP), <http://www.icir.org/kohler/dcp/>.
- [5] MAWI Working Group Traffic Archive, Packet traces from WIDE backbone, August 2003.
- [6] L. Ong, J. Yoakum, An Introduction to the Stream Control Transmission Protocol (SCTP), IETF RFC 3286, May 2002.
- [7] V. Padmanabhan, H. Balakrishnan, K. Sklower, E. Amir, R. Katz, Networking Using Direct Broadcast Satellite, In Proc. Workshop on Satellite-Based Information Systems (WOSBIS), November 1996.
- [8] N. Spring, D. Wetherall, A Protocol Independent Technique for Eliminating Redundant Network Traffic, In Proc. ACM SIGCOMM'00, August 2000.