Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: May 6, 2020

N. Kuhn
CNES
G. Fairhurst
University of Aberdeen
J. Border
Hughes Network Systems, LLC
E. Stephan
Orange

November 3, 2019

QUIC for SATCOM
draft-kuhn-quic-4-sat-02

Abstract

QUIC has been designed for use across Internet paths. Initial designs of QUIC have focussed on common deployment scenarios for web traffic and have not focussed on the performance when using a path with a large Bandwidth-Delay Product (BDP). A path can combine satellites network segment together with a wide variety of other network technologies (Ethernet, cable modems, WiFi, cellular, radio links, etc): this complicates the characteristics of the end-to-end path. One example of such a scenario occurs when a satellite communication (SATCOM) system is used to provide all or a part of the end-to-end path. If this is not addressed, the end-to-end quality of experience can be degraded.

This memo identifies the characteristics of a SATCOM link that impact the operation of the QUIC transport protocol and proposes current practice to ensure acceptable protocol performance.

Table of Contents

1. Introduction

The end-to-end performance of an application using an Internet path can be impacted by the Bandwidth-Delay Product (BDP) of the links and network devices forming the path. For instance, the page load time for a complex page can be much larger when the path includes a satellite link. A significant contribution to this reduced performance arises from the initialisation and design of transport mechanisms. QUIC's default congestion control is based on TCP NewReno [I-D.ietf-quic-recovery] and the recommended initial window is defined by [RFC6928]. Although QUIC's CC and recovery have been designed for use across Internet Paths, the initial design could not optimise for the wide diversity of path characteristics that can

occur. This document therefore considers the specific implications of paths with a significant BDP.

Satellite communications (SATCOM) systems have long been used to support point-to-point links and specialised networks. The predominate current use is as a link-layer for Internet Protocols. Typical example applications include: use as an access technology for remote locations, backup and rapid deployment of new services, transit networks and backhaul of various types of IP networks, and provision to mobile (maritime, aircraft, etc.). In most scenarios, the satellite IP network segment usually only forms one part of the end-to-end path. This means user traffic can experience a path that includes satellite link together with a wide variety of other network technologies (Ethernet, cable modems, WiFi, cellular, radio links, etc). Although a user can sometimes know the presence of the satellite service, a typical user does not deploy special software or applications because they expect a satellite network is being used. Often a user is unaware of the technologies underpinning the links forming the network path.

This memo identifies the characteristics of a SATCOM link that impact the operation of the QUIC transport protocol and proposes best current practice to ensure acceptable protocol performance.

2. Operating over a path with a large BDP

The characteristics of systems using Geosynchronous Earth Orbit (GEO) satellites differ from paths only using terrestrial links in their path characteristics:

- A large propagation delay of at least 250ms one-way delay;
- Employ radio resource management (often using techniques similar to cellular mobile or DOCSIS cable networks, but differing to accommodate the satellite propagation delay);
- Links can be highly asymmetric (in terms of capacity and one-way delay).

Many systems use the DVB-S2 specifications, where the key concept is to ensure both a good usage of the satellite resource and a Quasi Error Free (QEF) link. It consists in monitoring the link quality in real-time, with the help of known symbols sequences, included along regular packets, on which an estimation of the current signal-to-noise ratio can be done. Then, this estimation is send back to the transmitter that can adapt its coding rate and modulation order to best fit the actual transmission conditions.

It is common to consider the satellite network segment composed of a forward link and a return link. The two links can have different capacities and employ different technologies to carry the IP packets. On the forward link, the satellite gateway uses all the available bandwidth, possibly with several carriers, to communicate with the remote terminals. A carrier is a single Time-Division-Multiplexing where packets addressed to terminals are multiplexed. On the return link, the satellite resource is shared among the users. Two access methods can be distinguished: on-demand access or contention access. In the former, a terminal receives dedicated resources on its own to communicate with the gateway. In the latter, some resources are reserved for contention access, where several terminals can compete to obtain the resource. Dedicated access, which is more common in currently deployed systems, can be through a Demand Assigned Multiple Access (DAMA) mechanism, while contention access techniques are usually based on Slotted Aloha (SA) and its numerous derivatives. More information on satellite links characteristics can be found in [RFC2488][IJSCN17].

Beyond that, even for characteristics shared with terrestrial links, the impact on a satellite link could be more and can be amplified by the large RTT. For example, systems can exhibit a high loss-rate (e.g. mobile users or users behind a Wi-Fi link) which would impact loss recovery mechanisms and congestion control reaction to such loss events. The characteristics of a GEO SATCOM system impact the performance of congestion control:

o Transport initialization: the 3-way handshake takes a long time to complete, delaying the time at which actual data can be transmitted;

o Size of windows required: to fully exploit the bottleneck capacity, a high BDP will increase the number of in-flights packets;

o Reliability: packet loss detection and correction is slow (the performance of end-to-end retransmission is also impacted when using a high RTT path);

o Getting up to speed: the exponential increase of the data rate during slow start for a channel capacity probing is slowed down when the RTT is high;

o Asymetry : when the links are asymmetric, for various reasons, the sizing of the return link may induce modifications of the transport level acknowledgement traffic.

3. TCP Split Solution

High BDP networks commonly break the TCP end-to-end paradigm to adapt the transport protocol. Splitting TCP allows adaptations to this specific use-case and assessing the issues discussed in section Section 2. Satellite communications commonly deploy Performance Enhancement Proxy (PEP) for compression, caching and TCP acceleration services [RFC3135]. Their deployment can result in 50% page load time reduction in a SATCOM use-case [ICCRG100].

[NCT13] and [RFC3135] describe the main functions of SATCOM TCP split solutions. Shortly, for traffic originated at a gateway to an endpoint connected via a satellite terminal, the TCP split intercepts TCP SYN packets to act on behalf of the endpoint and adapt the data rate transmission to the SATCOM scenario. The split solution specifically tunes the TCP parameters to the context (latency, available capacity) of each link. When a PEP is used on each side of the satellite link, a protocol other than TCP, optimized for the satellite link, may be used. The tuning can be achieved using a priori information about the satellite system and/or by measuring the properties of the network segment that includes the satellite system.

One important advantage of a TCP split solution is that it does not require any end-to-end modifications and is independent for both client and server sides. That being said, this comes with a drawback: TCP splitters often are unable to track the most recent end-to-end improvements in protocol mechanisms (e.g., RACK, ECN, TCP Fast Open support) contributing to ossification of the transport system. The methods configured in the split proxy usually continue to be used until a split solution is finally updated. This can delay/negate the benefit of any end-to-end improvements.

4. Mechanisms that improve the performance of QUIC for SATCOM

4.1. Getting up to speed

The advantage of using QUIC is that it includes the TLS and TCP negociations that reduce the time at which the data can be transmitted. That being said, results of [IJSCN19] illustrate that it will take many RTTs for the congestion controller to increase the rate before it fills the bottleneck capacity. This dominates performance when a path has a large RTT (as in GEO SATCOM networks). There is an issue in QUIC getting up to speed in a SATCOM context.

The tuning of the initial window described in [I-D.irtf-iccrg-sallantin-initial-spreading] which has been shown to improve performance both for high BDP and more common BDP [CONEXT15][ICC16] could be a relevant solution. That being said,

such solution requires the usage of pacing to avoid important bursts of packets in the network that does not have a large BDP.

4.2. Maximum window

A large number of in-flight packets are prewired to fully exploit the bottleneck capacity, when there is a large BDP. Default values of maximum windows may not be suitable for a SATCOM context.

Such as presented in [PANRG105], only increasing the initial congestion window is not the only way that can improve QUIC performance in a SATCOM context: increasing maximum congestion windows can also result in much better performance. Other protocol mechanisms also need to be considered, such as flow control at the stream level in QUIC.

4.3. Reliability

Packet loss detection and loss repair take additional time on paths with a larger RTT. This increases the time that a server needs to react to a congestion event. Both can impact the user experience. This happens when a user uses a Wi-Fi link to access a SATCOM terminal. Although the benefits needed to weighed against the additional capacity in introducing end-to-end FEC and the potential to use link-local ARQ and/or link-adaptive FEC. End-to-end connections may not only suffer from losses in the Wi-Fi segment but also from congestion losses in the satellite operator ground segment. Using the mechanisms proposed in [I-D.ferrieux-hamchaoui-quic-lossbits], congestion losses have been identified on the ground segment.

Introducing network coding in QUIC such as proposed in [I-D.swett-nwcrg-coding-for-quic] and [I-D.roca-nwcrg-rlc-fec-scheme-for-quic] could help in recovering from the residual Wi-Fi or congestion losses. Another solution would be the usage of QUIC tunnels [I-D.schinazi-masque].

4.4. ACK ratio

Asymmetry in capacity (or in the way capacity is granted to a flow) can lead to cases where the throughput in one direction of communication is restricted by the acknowledgement traffic flowing in the opposite direction. The limitations of specific underlying networks could be in terms of the volume of acknowledgement traffic (limited return path capacity) or in the number of acknowledgement packets (e.g., when a radio-resource management system has to track channel usage) or both.

TCP Performance Implications of Network Path Asymmetry [RFC3449] describes a range of mechanisms that can mitigate the impact of path asymmetry. One simple method is to tell the remote endpoint to send compound acknowledgments less frequently. A rate of one ACK every RTT/4 can significantly reduce this traffic.

Many of these mitigations have been deployed in satellite systems, often as a mechanism within a PEP. Despite their benefits over paths with a high asymmetry of capacity, most mechanisms rely on being able to inspect and/or modify the transport layer header information of TCP ACK packets. This is not possible when the transport layer information is encrypted. The QUIC transport specification may evolve to allow the ACK Ratio to be adjusted.

5. Discussion

Many of the issues identified already exist for any encrypted transport service that uses a path that employs encryption at the IP layer. This includes endpoints that utilise IPsec at the network layer, or use VPN technology over the satellite network segment. These uses are unable to benefit from enhancement within the satellite network segment, and often the user is unaware of the presence of the satellite link on their path, except through observing the impact it has on the performance they experience.

One solution would be to provide PEP functions at the termination of the security association (e.g., in a VPN client). Another solution could be to fall-back to using TCP (possibly with TLS or similar methods being used on the transport payload). A final solution could be to deploy and maintain a bespoke protocol tailored to high BDP environments. In the future, we anticipate that fall-back will become less desirable, and methods that rely upon bespoke configurations or protocols will be unattractive. In parallel, new methods such as QUIC will become widely deployed. The opportunity therefore exists to ensure that the new generation of protocols offer acceptable performance over high BDP paths without requiring operating tuning or specific updates by users.

6. Acknowledgements

TBD

7. Contributors

TBD

8. IANA Considerations

TBD

9. Security Considerations

This document does not propose changes to the security functions provided by the QUIC protocol. QUIC uses TLS encryption to protect the transport header and its payload. Security is considered in the "Security Considerations" of cited IETF documents.

10. References

10.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

10.2. Informative References

[CONEXT15]
Li, Q., Dong, M., and P. Godfrey, "Halfback: Running Short Flows Quickly and Safely", ACM CoNEXT , 2015.

[I-D.ferrieux-hamchaoui-quic-lossbits]
Ferrieux, A. and I. Hamchaoui, "The QUIC Loss Bits", draft-ferrieux-hamchaoui-quic-lossbits-00 (work in progress), April 2019.

[I-D.ietf-quic-recovery]
Iyengar, J. and I. Swett, "QUIC Loss Detection and Congestion Control", draft-ietf-quic-recovery-23 (work in progress), September 2019.

[I-D.ietf-quic-tls]
Thomson, M. and S. Turner, "Using TLS to Secure QUIC", draft-ietf-quic-tls-23 (work in progress), September 2019.

[I-D.ietf-quic-transport]
Iyengar, J. and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport", draft-ietf-quic-transport-23 (work in progress), September 2019.

[I-D.ietf-tls-ticketrequests]
Pauly, T., Schinazi, D., and C. Wood, "TLS Ticket Requests", draft-ietf-tls-ticketrequests-03 (work in progress), October 2019.

[I-D.irtf-iccrg-sallantin-initial-spreading]
Sallantin, R., Baudoin, C., Arnal, F., Dubois, E., Chaput, E., and A. Beylot, "Safe increase of the TCP's Initial Window Using Initial Spreading", draft-irtf-iccrg-sallantin-initial-spreading-00 (work in progress), January 2014.

[I-D.roca-nwcrg-rlc-fec-scheme-for-quic]
Roca, V., Swett, I., and M. Montpetit, "Sliding Window Random Linear Code (RLC) Forward Erasure Correction (FEC) Schemes for QUIC", draft-roca-nwcrg-rlc-fec-scheme-for-quic-01 (work in progress), February 2019.

[I-D.schinazi-masque]
Schinazi, D., "The MASQUE Protocol", draft-schinazi-masque-01 (work in progress), July 2019.

[I-D.swett-nwcrg-coding-for-quic]
Swett, I., Montpetit, M., Roca, V., and F. Michel, "Coding for QUIC", draft-swett-nwcrg-coding-for-quic-03 (work in progress), July 2019.

[ICC16] Sallantin, R., Baudoin, C., Chaput, E., Arnal, F., Dubois, E., and A-L. Beylot, "Reducing web latency through TCP IW: Be smart", IEEE ICC , 2016.

[ICCRG100]
Kuhn, N., "MPTCP and BBR performance over Internet satellite paths", IETF ICCRG 100, 2017.

[IJSCN17] Ahmed, T., Dubois, E., Dupe, JB., Ferrus, R., Gelard, P., and N. Kuhn, "Software-defined satellite cloud RAN", International Journal of Satellite Communications and Networking , 2017.

[IJSCN19] Thomas, L., Dubois, E., Kuhn, N., and E. Lochin, "Google QUIC performance over a public SATCOM access", International Journal of Satellite Communications and Networking , 2019.

[NCT13] Pirovano, A. and F. Garcia, "A new survey on improving TCP performances over geostationary satellite link", Network and Communication Technologies , 2013.

[PANRG105]
Kuhn, N., Stephan, E., Border, J., and G. Fairhurst, "QUIC Over In-sequence Paths with Different Characteristics", IRTF PANRG 105, 2019.

[RFC2488] Allman, M., Glover, D., and L. Sanchez, "Enhancing TCP Over Satellite Channels using Standard Mechanisms", BCP 28, RFC 2488, DOI 10.17487/RFC2488, January 1999, <https://www.rfc-editor.org/info/rfc2488>.

[RFC3135] Border, J., Kojo, M., Griner, J., Montenegro, G., and Z. Shelby, "Performance Enhancing Proxies Intended to Mitigate Link-Related Degradations", RFC 3135, DOI 10.17487/RFC3135, June 2001, <https://www.rfc-editor.org/info/rfc3135>.

[RFC3449] Balakrishnan, H., Padmanabhan, V., Fairhurst, G., and M. Sooriyabandara, "TCP Performance Implications of Network Path Asymmetry", BCP 69, RFC 3449, DOI 10.17487/RFC3449, December 2002, <https://www.rfc-editor.org/info/rfc3449>.

[RFC6349] Constantine, B., Forget, G., Geib, R., and R. Schrage, "Framework for TCP Throughput Testing", RFC 6349, DOI 10.17487/RFC6349, August 2011, <https://www.rfc-editor.org/info/rfc6349>.

[RFC6928] Chu, J., Dukkipati, N., Cheng, Y., and M. Mathis, "Increasing TCP's Initial Window", RFC 6928, DOI 10.17487/RFC6928, April 2013, <https://www.rfc-editor.org/info/rfc6928>.

[RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <https://www.rfc-editor.org/info/rfc8446>.

Authors' Addresses

Nicolas Kuhn
CNES

Email: nicolas.kuhn@cnes.fr

Godred Fairhurst
University of Aberdeen

Email: gorry@erg.abdn.ac.uk

John Border
Hughes Network Systems, LLC

Email: border@hns.com

Emile Stephan
Orange

Email: emile.stephan@orange.com