Transport Area Working Group
Internet-Draft
Intended status: Experimental
Expires: September 12, 2019

M. Amend
Deutsche Telekom
A. Brunstrom
A. Kassler
Karlstad University
V. Rakocevic
City University of London

March 11, 2019

DCCP Extensions for Multipath Operation with Multiple Addresses
draft-amend-tsvwg-multipath-dccp-01

Abstract

DCCP communication is currently restricted to a single path per connection, yet multiple paths often exist between peers. The simultaneous use of these multiple paths for a DCCP session could improve resource usage within the network and, thus, improve user experience through higher throughput and improved resilience to network failure.

Multipath DCCP provides the ability to simultaneously use multiple paths between peers. This document presents a set of extensions to traditional DCCP to support multipath operation. The protocol offers the same type of service to applications as DCCP and it provides the components necessary to establish and use multiple DCCP flows across potentially disjoint paths.

Copyright Notice

Table of Contents

1. Introduction

Multipath DCCP (MP-DCCP) is a set of extensions to regular DCCP [RFC4340], which enables a transport connection to operate across multiple paths simultaneously. DCCP multipath operations is suggested in the context of ongoing 3GPP work on 5G multi-access solutions [draft-amend-tsvwg-multipath-framework-mpdccp] and for hybrid access networks [draft-lhwxz-hybrid-access-network-architecture, draft-muley-network-based-bonding-hybrid-access]. It can be applied for load-balancing, seamless session handover and aggregation purposes (referred to as steering, switching and splitting in 3GPP terminology [3GPP, TR 23.793]).

This document presents the protocol changes required to add multipath capability to DCCP; specifically, those for signaling and setting up

multiple paths ("subflows"), managing these subflows, reassembly of data, and termination of sessions.

1.1. Multipath DCCP in the Networking Stack

MP-DCCP operates at the transport layer and aims to be transparent to both higher and lower layers. It is a set of additional features on top of standard DCCP; Figure 1 illustrates this layering. MP-DCCP is designed to be used by applications in the same way as DCCP with no changes.

```
                             +-------------------------------+
                             |           Application         |
+---------------+            +-------------------------------+
|  Application  |            |            MP-DCCP            |
+---------------+            + - - - - - - - + - - - - - - - +
|      DCCP     |            |Subflow (DCCP) |Subflow (DCCP) |
+---------------+            +-------------------------------+
|      IP       |            |      IP       |      IP       |
+---------------+            +-------------------------------+
```

Figure 1: Comparison of Standard DCCP and MP-DCCP Protocol Stacks

1.2. Terminology

[Tbd], could be similar to [RFC6824]

1.3. MP-DCCP Concept

```
           Host A                                Host B
------------------------             ------------------------
Address A1    Address A2             Address B1    Address B2
----------    ----------             ----------    ----------
  |             |                      |             |
  |         (DCCP flow setup)          |             |
  |----------------------------------->|             |
  |<-----------------------------------|             |
  |             |                      |             |
  |             |  (DCCP flow setup)   |             |
  |             |--------------------->|             |
  |             |<---------------------|             |
  | merge individual DCCP flows to one multipath connection
  |             |                      |             |
```

Figure 2: Example MP-DCCP Usage Scenario

1.4. Differences from Multipath TCP

Multipath DCCP is similar to Multipath TCP [RFC6824], in that it extends the related basic DCCP transport protocol [RFC4340] with multipath capabilities in the same way as Multipath TCP extends TCP [RFC793]. However, mainly dominated by the basic protocols TCP and DCPP, the transport characteristics are different.

Table 1 compares the protocol characteristics of TCP and DCCP, which are by nature inherited by their respective multipath extensions. A major difference lies in the delivery of payload, which is for TCP an exact copy of the generated byte-stream. DCCP behaves contrary and does not guarantee to transmit any payload nor the order of delivery. Since this is mainly affecting the receiving endpoint of a TCP or DCCP communication, many similarities on sender side can be stated. Both transport protocols share the 3-way initiation of a communication and both exploit a congestion control to adapt to path characteristics.

Feature	TCP	DCCP
Full-Duplex	yes	yes
Connection-Oriented	yes	yes
Header option space	40 bytes	< 1008 bytes or PMTU
Data transfer	reliable	unreliable
Packet-loss handling	re-transmission	report only
Ordered data delivery	yes	no
Sequence numbers	one per byte	one per PDU
Flow control	yes	no
Congestion control	yes	yes
ECN support	yes	yes

Selective ACK	yes	depends on congestion control
Fix message boundaries	no	yes
Path MTU discovery	yes	yes
Fragmentation	yes	no
SYN flood protection	yes	no
Half-open connections	yes	no

Table 1: TCP and DCCP protocol comparison

Consequently, the multipath features, shown in Table 2, are the same for support of volatile paths, session handover and path aggregation capabilities. All of them profit by the existence of congestion control.

Feature	MP-TCP	MP-DCCP
Volatile paths	yes	yes
Robust session establishment	no	yes
Data reassembly	yes	optional
Expandability	limited by TCP header	flexible
Session handover	yes	yes
Path aggregation	yes	yes

Table 2: MPTCP and MP-DCCP protocol comparison

Therefore the sender logic is not much different between MP-DCCP and MPTCP, even if the multipath session initiation differs. MP-DCCP inherits a robust session establishment feature, which guarantees communication establishment if at least one functional path is available. MP-TCP relies on an initial path, which has to work; otherwise no communication can be established.

The receiver side for MP-DCCP has to deal with the unreliable transport character of DCCP and a possible re-assembly of the data stream. In practice, it is assumed that some sort of re-assembly has to be applied, even if DCCP and the order of delivery is unreliable by nature. Such re-assembly mechanisms have to account for the fact that packet loss may occur for any of the DCCP subflows. Another issue is the packet reordering introduced when a DCCP communication is split across paths with disjoint latencies. In theory, applications using DCCP certainly have to deal with packet reordering, since DCCP has no mechanisms to prevent it. However, in practice, without any multipath extension, packet reordering can be assumed to be very limited. Therefore most services on top of DCCP are not expecting massive packet reordering and degrades their performance if it happens anyway.

The receiving process for MP-TCP is on the other hand a simple "just wait" approach, since TCP guarantees reliable delivery.

1.5. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2. Operation Overview

[Tbd], could be similar to [RFC6824]

3. MP-DCCP Protocol

[Tbd], could be similar to [RFC6824]

[Tbd] On top it requires particular considerations for:

o The minimum PMTU of the individual paths must be selected to announce to the application. Changes of individual path PMTUs must be re-announced to the application if they are lower than the current announced PMTU.

- o Overall sequencing for optional reassembly procedure

- o Congestion control

- o Robust MP-DCCP session establishment (no dependency on an initial path setup)

4. Security Considerations

[Tbd]

5. Interactions with Middleboxes

[Tbd], should mention standardized technologies like [RFC5597] or [RFC6773] and U-DCCP [draft-amend-tsvwg-dccp-udp-header-conversion]

6. Acknowledgments

1. Notes

This document is inspired by Multipath TCP [RFC6824] and some text passages for the -00 version of the draft are copied almost unmodified.

7. IANA Considerations

[Tbd], must include options for:

- o handshaking procedure to indicate MP support

- o handshaking procedure to indicate JOINING of an existing MP connection

- o signaling of new or changed addresses

- o setting handover or aggregation mode

- o setting reordering on/off

should include options carrying:

- o overall sequence number for restoring purposes

- o sender time measurements for restoring purposes

- o scheduler preferences

- o reordering preferences

8. Informative References

[I-D.amend-tsvwg-multipath-framework-mpdccp]
Amend, M., Brunstrom, A., Kassler, A., and V. Rakocevic, "IP compatible multipath framework for heterogeneous access networks", draft-amend-tsvwg-multipath-framework-mpdccp-00 (work in progress), March 2019.

[I-D.lhwxz-hybrid-access-network-architecture]
Leymann, N., Heidemann, C., Wasserman, M., Xue, L., and M. Zhang, "Hybrid Access Network Architecture", draft-lhwxz-hybrid-access-network-architecture-02 (work in progress), January 2015.

[I-D.muley-network-based-bonding-hybrid-access]
Muley, P., Henderickx, W., Geng, L., Liu, H., Cardullo, L., Newton, J., Seo, S., Draznin, S., and B. Patil, "Network based Bonding solution for Hybrid Access", draft-muley-network-based-bonding-hybrid-access-03 (work in progress), October 2018.

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <https://www.rfc-editor.org/info/rfc793>.

[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>.

[RFC4340] Kohler, E., Handley, M., and S. Floyd, "Datagram Congestion Control Protocol (DCCP)", RFC 4340, DOI 10.17487/RFC4340, March 2006, <https://www.rfc-editor.org/info/rfc4340>.

[RFC5597] Denis-Courmont, R., "Network Address Translation (NAT) Behavioral Requirements for the Datagram Congestion Control Protocol", BCP 150, RFC 5597, DOI 10.17487/RFC5597, September 2009, <https://www.rfc-editor.org/info/rfc5597>.

[RFC6773] Phelan, T., Fairhurst, G., and C. Perkins, "DCCP-UDP: A Datagram Congestion Control Protocol UDP Encapsulation for NAT Traversal", RFC 6773, DOI 10.17487/RFC6773, November 2012, <https://www.rfc-editor.org/info/rfc6773>.

[RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, DOI 10.17487/RFC6824, January 2013, <https://www.rfc-editor.org/info/rfc6824>.

Authors' Addresses

Markus Amend
Deutsche Telekom
T-Online-Allee 1
Darmstadt
Germany

Email: Markus.Amend@telekom.de

Anna Brunstrom
Karlstad University
Universitetsgatan 2
651 88 Karlstad
Sweden

Email: anna.brunstrom@kau.se

Andreas Kassler
Karlstad University
Universitetsgatan 2
651 88 Karlstad
Sweden

Email: andreas.kassler@kau.se

Veselin Rakocevic
City University of London
Northampton Square
London
United Kingdom

Email: veselin.rakocevic.1@city.ac.uk