

AVT Core Working Group
Internet-Draft
Intended status: Experimental
Expires: January 9, 2012

V. Singh
T. Karkkainen
J. Ott
S. Ahsan
Aalto University
L. Eggert
Nokia
July 8, 2011

Multipath RTP (MPRTP)
draft-singh-avtcore-mprtp-02

Abstract

The Real-time Transport Protocol (RTP) is used to deliver real-time content and, along with the RTP Control Protocol (RTCP), forms the control channel between the sender and receiver. However, RTP and RTCP assume a single delivery path between the sender and receiver and make decisions based on the measured characteristics of this single path. Increasingly, endpoints are becoming multi-homed, which means that they are connected via multiple Internet paths. Network utilization can be improved when endpoints use multiple parallel paths for communication. The resulting increase in reliability and throughput can also enhance the user experience. This document extends the Real-time Transport Protocol (RTP) so that a single session can take advantage of the availability of multiple paths between two endpoints.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 9, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Requirements Language	4
1.2. Terminology	4
1.3. Use-cases	5
2. Goals	5
2.1. Functional goals	5
2.2. Compatibility goals	6
3. RTP Topologies	6
4. MP RTP Architecture	6
4.1. Relationship of MP RTP with Session Signaling	8
5. Example Media Flow Diagrams	8
5.1. Streaming use-case	8
5.2. Conversational use-case	9
5.3. Challenges with Multipath Interface Discovery	10
6. MP RTP Functional Blocks	10
7. Available Mechanisms within the Functional Blocks	11
7.1. Session Setup	11
7.2. Expanding RTP	11
7.3. Adding New Interfaces	11
7.4. Expanding RTCP	12
7.5. Checking and Failure Handling	12
8. MP RTP Protocol	12
8.1. Overview	13
8.1.1. Subflow or Interface advertisement	14
8.1.2. Path selection	14
8.1.3. Opening subflows	15
8.2. RTP Transmission	15
8.3. Playout Considerations at the Receiver	15
8.4. Subflow-specific RTCP Statistics and RTCP Aggregation	16
8.5. RTCP Transmission	16
9. Packet Formats	16
9.1. RTP Header Extension for MP RTP	16

9.1.1. MPRTP RTP Extension for a Subflow	18
9.1.2. MPRTP RTP Extension for Connectivity Checks	19
9.1.3. MPRTP RTP Extension for Keep-alive Packets	19
9.2. RTCP Extension for MPRTP (MPRTCP)	19
9.2.1. MPRTCP Extension for Subflow Reporting	20
9.2.1.1. MPRTCP for Subflow-specific SR, RR and XR	22
9.3. MPRTCP Extension for Interface advertisement	24
9.3.1. Interface Advertisement block	25
10. RTCP Timing reconsiderations for MPRTCP	26
11. SDP Considerations	26
11.1. Signaling MPRTP capability in SDP	26
11.2. Signaling MPRTP Header Extension in SDP	27
11.3. MPRTP using preloaded interfaces from ICE (out-band signaling)	27
11.3.1. "interface" attribute	27
11.3.2. Example	28
11.4. Increased Throughput	29
11.5. Offer/Answer Examples	29
11.5.1. In-band signaling	29
11.5.2. Out-band signaling	30
12. IANA Considerations	30
13. Security Considerations	31
14. Acknowledgements	31
15. References	31
15.1. Normative References	31
15.2. Informative References	32
Authors' Addresses	33

1. Introduction

Multi-homed endpoints are becoming common in today's Internet, e.g., devices that support multiple wireless access technologies such as 3G and Wireless LAN. This means that there is often more than one network path available between two endpoints. Transport protocols, such as RTP, have not been designed to take advantage of the availability of multiple concurrent paths and therefore cannot benefit from the increased capacity and reliability that can be achieved by pooling their respective capacities.

Multipath RTP (MPRTP) is an OPTIONAL extension to RTP [1] that allows splitting a single RTP stream into multiple subflows that are transmitted over different paths. In effect, this pools the resource capacity of multiple paths. Multipath RTCP (MPRTCP) is an extension to RTCP, it is used along with MPRTP to report per-path sender and receiver characteristics.

Other IETF transport protocols that are capable of using multiple paths include SCTP [10], MPTCP [11] and SHIM6 [12]. However, these protocols are not suitable for realtime communications.

1.1. 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 [2].

1.2. Terminology

- o Endpoint: host either initiating or terminating an RTP connection.
- o Interface: logical or physical component that is capable of acquiring a unique IP address.
- o Path: sequence of links between a sender and a receiver. Typically, defined by a set of source and destination addresses.
- o Subflow: flow of RTP packets along a specific path, i.e., a subset of the packets belonging to an RTP stream. The combination of all RTP subflows forms the complete RTP stream. Typically, a subflow would map to a unique path, i.e., each combination of IP addresses and port pairs (5-tuple) is a unique subflow.

1.3. Use-cases

The primary use-case for MP RTP is transporting high bit-rate streaming multimedia content between endpoints, where at least one is multi-homed. Such endpoints could be residential IPTV devices that connect to the Internet through two different Internet service providers (ISPs), or mobile devices that connect to the Internet through 3G and WLAN interfaces. By allowing RTP to use multiple paths for transmission, the following gains can be achieved:

- o Higher quality: Pooling the resource capacity of multiple Internet paths allows higher bit-rate and higher quality codecs to be used. From the application perspective, the available bandwidth between the two endpoints increases.
- o Load balancing: Transmitting one RTP stream over multiple paths can reduce the bandwidth usage, compared to transmitting the same stream along a single path. This reduces the impact on other traffic.
- o Fault tolerance: When multiple paths are used in conjunction with redundancy mechanisms (FEC, re-transmissions, etc.), outages on one path have less impact on the overall perceived quality of the stream.

A secondary use-case for MP RTP is transporting Voice over IP (VoIP) calls to a device with multiple interfaces. Again, such an endpoint could be a mobile device with multiple wireless interfaces. In this case, little is to be gained from resource pooling, i.e., higher capacity or load balancing, because a single path should be easily capable of handling the required load. However, using multiple concurrent subflows can improve fault tolerance, because traffic can shift between the subflows when path outages occur. This results in very fast transport-layer handovers that do not require support from signaling.

2. Goals

This section outlines the basic goals that multipath RTP aims to meet. These are broadly classified as Functional goals and Compatibility goals.

2.1. Functional goals

Allow unicast RTP session to be split into multiple subflows in order to be carried over multiple paths. This may prove beneficial in case of video streaming.

- o Increased Throughput: Cumulative capacity of the two paths may meet the requirements of the multimedia session. Therefore, MP RTP MUST support concurrent use of the multiple paths.
- o Improved Reliability: MP RTP SHOULD be able to send redundant packets or re-transmit packets along any available path to increase reliability.

The protocol SHOULD be able to open new subflows for an existing session when new paths appear and MUST be able to close subflows when paths disappear.

2.2. Compatibility goals

MP RTP MUST be backwards compatible; an MP RTP stream needs to fall back to be compatible with legacy RTP stacks if MP RTP support is not successfully negotiated.

- o Application Compatibility: MP RTP service model MUST be backwards compatible with existing RTP applications, i.e., an MP RTP stack MUST be able to work with legacy RTP applications and not require changes to them. Therefore, the basic RTP APIs MUST remain unchanged, but an MP RTP stack MAY provide extended APIs so that the application can configure any additional features provided by the MP RTP stack.
- o Network Compatibility: individual RTP subflows MUST themselves be well-formed RTP flows, so that they are able to traverse NATs and firewalls. This MUST be the case even when interfaces appear after session initiation. Interactive Connectivity Establishment (ICE) [3] MAY be used for discovering new interfaces or performing connectivity checks.

3. RTP Topologies

RFC 5117 [13] describes a number of scenarios using mixers and translators in single-party (point-to-point), and multi-party (point-to-multipoint) scenarios. RFC 3550 [1] (Section 2.3 and 7.x) discuss in detail the impact of mixers and translators on RTP and RTCP packets. MP RTP assumes that if a mixer or translator exists in the network, then either all of the multiple paths or none of the multiple paths go via this component.

4. MP RTP Architecture

In a typical scenario, an RTP session uses a single path. In an

MPRTP scenario, an RTP session uses multiple subflows that each use a different path. Figure 1 shows the difference.




Figure 1: Comparison between traditional RTP streaming and MPRTP




Figure 2: MPRTP Architecture

Figure 2 illustrates the differences between the standard RTP stack and the MPRTP stack. MPRTP receives a normal RTP session from the application and splits it into multiple RTP subflows. Each subflow is then sent along a different path to the receiver. To the network, each subflow appears as an independent, well-formed RTP flow. At the receiver, the subflows are combined to recreate the original RTP session. The MPRTP layer performs the following functions:

- o Path Management: The layer is aware of alternate paths to the other host, which may, for example, be the peer's multiple interfaces. So that it is able to send differently marked packets along separate paths. MPRTP also selects interfaces to send and receive data. Furthermore, it manages the port and IP address pair bindings for each subflow.

- o Packet Scheduling: the layer splits a single RTP flow into multiple subflows and sends them across multiple interfaces (paths). The splitting MAY BE done using different path characteristics.
- o Subflow recombination: the layer creates the original stream by recombining the independent subflows. Therefore, the multipath subflows appear as a single RTP stream to applications.

4.1. Relationship of MPRTTP with Session Signaling

Session signaling (e.g., SIP [14], RTSP [15]) SHOULD be done over a failover-capable or multipath-capable transport for e.g., SCTP [10] or MPTCP [11] instead of TCP or UDP.

5. Example Media Flow Diagrams

There may be many complex technical scenarios for MPRTTP, however, this memo only considers the following two scenarios: 1) a unidirectional media flow that represents the streaming use-case, and 2) a bidirectional media flow that represents a conversational use-case.

5.1. Streaming use-case

In the unidirectional scenario, the receiver (client) initiates a multimedia session with the sender (server). The receiver or the sender may have multiple interfaces and both endpoints are MPRTTP-capable. Figure 3 shows this scenario. In this case, host A has multiple interfaces. Host B performs connectivity checks on host A's multiple interfaces. If the interfaces are reachable, then host B streams multimedia data along multiple paths to host A. Moreover, host B also sends RTCP Sender Reports (SR) for each subflow and host A responds with a standard RTCP Receiver Report (RR) for the overall session and receiver statistics for each subflow. Host B distributes the packets across the subflows based on the individually measured path characteristics.

Alternatively, to reduce media startup time, host B may start streaming multimedia data to host A's initiating interface and then perform connectivity checks for the other interfaces. This method of updating a single path session to a multipath session is called "multipath session upgrade".

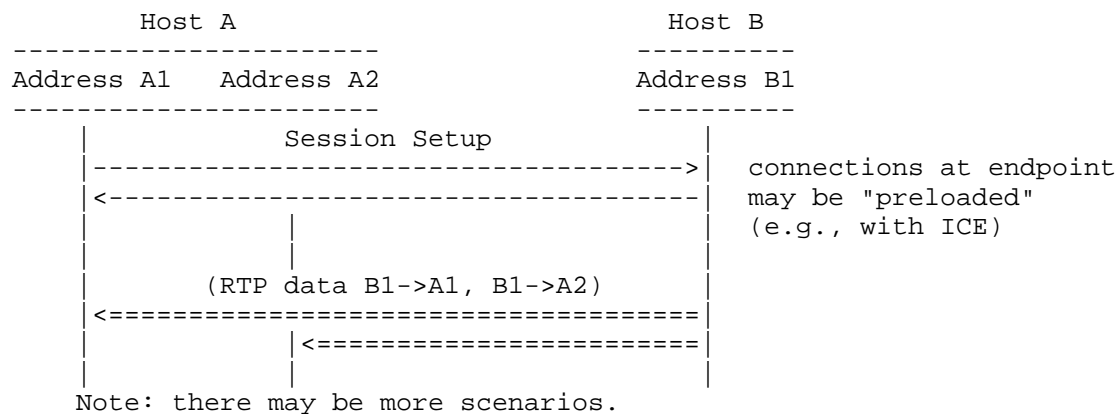


Figure 3: Unidirectional media flow

5.2. Conversational use-case

In the bidirectional scenario, multimedia data flows in both directions. The two hosts exchange their lists of interfaces with each other and perform connectivity checks. Communication begins after each host finds suitable address, port pairs. Interfaces that receive data send back RTCP receiver statistics for that path (based on the 5-tuple). The hosts balance their multimedia stream across multiple paths based on the per path reception statistics and its own volume of traffic. Figure 4 describes an example of a bidirectional flow.

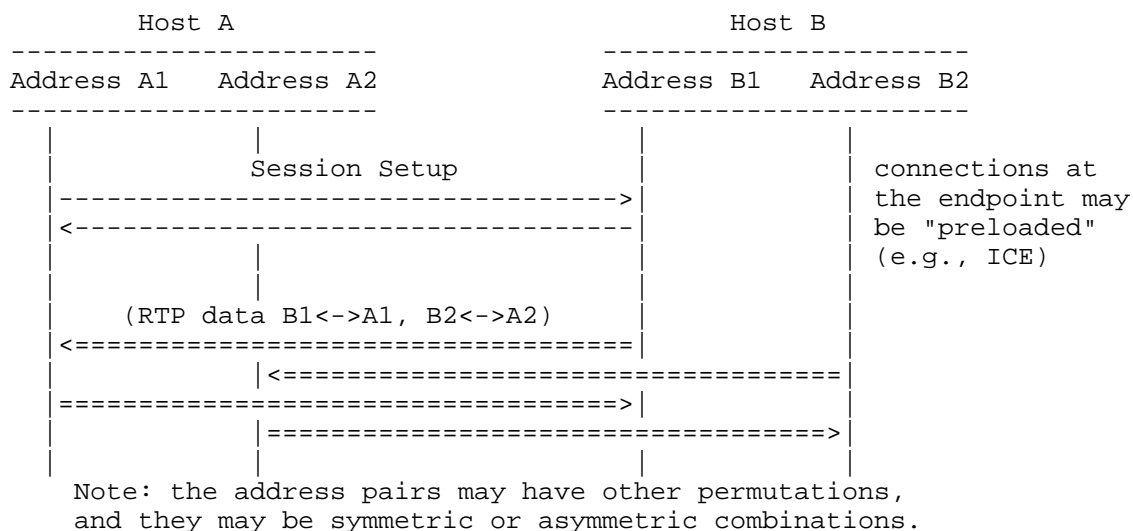


Figure 4: Bidirectional flow

5.3. Challenges with Multipath Interface Discovery

For some applications, where the user expects immediate playback, e.g., High Definition Media Streaming or IPTV, it may not be possible to perform connectivity checks within the given time bound. In these cases, connectivity checks MAY need to be done ahead of time.

[Open Issue: ICE or any other system would have to be aware of the endpoint's interfaces ahead of time].

6. MPRTTP Functional Blocks

This section describes some of the functional blocks needed for MPRTTP. We then investigate each block and consider available mechanisms in the next section.

1. **Session Setup:** Multipath session setup is an upgrade or add-on to a typical RTP session. Interfaces may appear or disappear at anytime during the session. To preserve backward compatibility with legacy applications, a multipath session MUST look like a bundle of individual RTP sessions.
2. **Expanding RTP:** For a multipath session, each subflow MUST look like an independent RTP flow, so that individual RTCP messages can be generated per subflow. Furthermore, MPRTTP splits the single multimedia stream into multiple subflows based on path characteristics (e.g. RTT, loss-rate, receiver rate, bandwidth-delay product etc.) and dynamically adjusts the load on each link.
3. **Adding Interfaces:** Interfaces on the host need to be regularly discovered and signaled. This can be done at session setup and/or during the session. When discovering and receiving new interfaces, the MPRTTP layer needs to select address and port pairs.
4. **Expanding RTCP:** MPRTTP MUST recombine RTCP reports from each path to re-create a single RTCP message to maintain backward compatibility with legacy applications.
5. **Maintenance and Failure Handling:** In a multi-homed endpoint interfaces may appear and disappear. If this happens at the sender, it has to re-adjust the load on the available links. On the other hand, if this occurs on the receiver, then the multimedia data transmitted by the sender to those interfaces is

lost. This data may be re-transmitted along a different path i.e., to a different interface on the receiver. Furthermore, the receiver has to explicitly signal the disappearance of an interface, or the sender has to detect it. [Open Issue: What happens if the interface that setup the session disappears? does the control channel also failover? re-start the session?]

6. Teardown: The MP RTP layer releases the occupied ports on the interfaces.

7. Available Mechanisms within the Functional Blocks

This section discusses some of the possible alternatives for each functional block mentioned in the previous section.

7.1. Session Setup

MP RTP session can be set up in many possible ways e.g., during handshake, or upgraded mid-session. The capability exchange may be done using out-of-band signaling (e.g., SDP [16] in SIP [14], RTSP [15]) or in-band signaling (e.g., RTP/RTCP header extension). Furthermore, ICE [3] may be used for discovering and performing connectivity checks during session setup.

7.2. Expanding RTP

RTCP [1] is generated per media session. However, with MP RTP, the media sender spreads the RTP load across several interfaces. The media sender SHOULD make the path selection, load balancing and fault tolerance decisions based on the characteristics of each path. Therefore, apart from normal RTP sequence numbers defined in [1], the MP RTP sender MUST add subflow-specific sequence numbers to help calculate fractional losses, jitter, RTT, playout time, etc., for each path and a subflow identifier to associate the characteristics with a path. The RTP header extension for MP RTP is shown in Section 9).

7.3. Adding New Interfaces

When interfaces appear and disappear mid-session, ICE [3] may be used for discovering interfaces and performing connectivity checks. However, MP RTP may require a capability re-negotiation (using SDP) to include all these new interfaces. This method is referred to as out-of-band multipath advertisement.

Alternatively, when new interfaces appear, the interface advertisements may be done in-band using RTP/RTCP extensions. The

endpoints perform connectivity checks (see Figure 5 for more details). If the connectivity packets are received by the peers, then multimedia data can flow between the new address, port pairs.

7.4. Expanding RTCP

To provide accurate per path information an MPRTP endpoint MUST send (SR/RR) report for each unique subflow along with the overall session RTCP report. Therefore, the additional subflow reporting affects the RTCP bandwidth and the RTCP reporting interval for each subflow. RTCP report scheduling for each subflow may cause a problem for RTCP recombination and reconstruction in cases when 1) RTCP for a subflow is lost, and 2) RTCP for a subflow arrives later than the other subflows. (There may be other cases as well.)


The sender distributes the media across different paths using the per path RTCP reports. However, this document doesn't cover algorithms for congestion control or load balancing.

7.5. Checking and Failure Handling

[Note: If the original interface that setup the session disappears then does the session signaling failover to another interface? Can we recommend that SIP/RTSP be run over MPTCP, SCTP].

8. MPRTP Protocol

To enable a quick start of a multimedia session, a multipath session MUST be upgraded from a single path session. Therefore, no explicit changes are needed in multimedia session setup and the session can be setup as before.



Key:
 | Interface
 ---> Signaling Protocol
 <=== RTP Packets
 - -> RTCP Packet

Figure 5: MPRTTP New Interface

8.1. Overview

The bullet points explain the different steps shown in Figure 5 for upgrading a standard single path multimedia session to multipath session.

(1) The first two interactions between the hosts represents the standard session setup. This may be SIP or RTSP.

(2) Following the setup, like in a conventional RTP scenario, host B using interface B1 starts to stream data to host A at interface A1.

(3) Host B is an MPRTTP-capable media sender and becomes aware of another interface B2.

(4) Host B advertises the multiple interface addresses using an RTCP header extensions.

(5) Host A is an MPRTTP-capable media receiver and becomes aware of another interface A2. It advertises the multiple interface addresses using an RTCP extension.

Side note, even if an MPRTTP-capable host has only one interface, it SHOULD respond to the advertisement with its single interface.

(6) Each host receives information about the additional interfaces and performs the connectivity tests (not shown in figure). If the paths are reachable then the host starts to stream the multimedia content using the additional paths.

8.1.1. Subflow or Interface advertisement

To advertise the multiple interfaces, an MPRTTP-capable endpoint MUST add the MPRTTP Interface Advertisement defined in Figure 13 with the RTCP Sender Report (SR). Each unique address is encapsulated in an Interface Advertisement block and contains the IP address, RTP and RTCP port addresses. The Interface Advertisement blocks are ordered based on a decreasing priority level. On receiving the MPRTTP Interface Advertisement, an MPRTTP-capable receiver MUST respond with its own set of interfaces.

If the sender and receiver have only one interface, then the endpoints MUST respond with the default IP, RTP port and RTCP port addresses. If an endpoint receives an RTCP report without the MPRTTP Interface Advertisement, then the endpoint MUST assume that the other endpoint is not MPRTTP capable.

8.1.2. Path selection

After MPRTTP support has been discovered and interface advertisements have been exchanged, the sender MUST initiate connectivity checks to determine which interface pairs offer valid paths between the sender and the receiver. Each combination of IP addresses and port pairs (5-tuple) is a unique subflow. An endpoint MUST associate a Subflow ID to each unique subflow.

To initiate a connectivity check, the endpoints send an RTP packet using the appropriate MPRTTP extension header (See Figure 7), associated Subflow ID and no RTP payload. The receiving endpoint replies to each connectivity check with an RTCP packet with the appropriate packet type (See Figure 10) and Subflow ID. After the endpoint receives the reply, the path is considered a valid candidate for sending data. An endpoint MAY choose to do any number of

connectivity checks for any interface pairs at any point in a session.

[Open Issue: How should the endpoint adjust the RTCP Reporting interval/schedule the RTCP packet on receiving a connectivity check containing a new Subflow ID? Editor: One option is send immediately as defined in [4]. Another option is the RTCP timing defined in [17].]

8.1.3. Opening subflows

The sender MAY open any number of subflows from the set of candidate subflows after performing connectivity checks. To use the subflow, the sender simply starts sending the RTP packets with an MPRTTP extension shown in Figure 8. The MPRTTP extension carries a mapping of a subflow packet to the aggregate flow. Namely, sequence numbers and timestamps associated with the subflow.

An endpoint MAY use all or a subset of candidate subflows for sending media packets. To avoid redoing the connectivity checks the endpoint MAY send keep-alive MPRTTP packets (see Section 9.1.3) to the passive subflows to keep the NAT bindings alive.

[Open Issue: How to differentiate between Passive and Active connections? Editor: Active paths get "regular flow" of media packets while passive paths are for failover of active paths.]

[Open Issue: How to keep a passive connection alive, if not actively used? Alternatively, what is the maximum timeout? Editor: keep-alive for ICE/NAT bindings should not be less than 15 seconds [3].]

8.2. RTP Transmission

The MPRTTP layer SHOULD associate an RTP packet with a subflow based on a scheduling strategy. The scheduling strategy may either choose to augment the paths to create higher throughput or use the alternate paths for enhancing resilience or error-repair. Due to the changes in path characteristics, an MPRTTP sender can change its scheduling strategy during an ongoing session. The MPRTTP sender MUST also populate the subflow specific fields described in the MPRTTP extension header (see Section 9.1.1).

8.3. Playout Considerations at the Receiver

A media receiver, irrespective of MPRTTP support or not, should be able to playback the media stream because the received RTP packets are compliant to [1], i.e., a non-MPRTTP receiver will ignore the MPRTTP header and still be able to playback the RTP packets. However,

the variation of jitter and loss per path may affect proper playout. The receiver can compensate for the jitter by modifying the playout delay (i.e., by calculating skew across all paths) of the received RTP packets.

8.4. Subflow-specific RTCP Statistics and RTCP Aggregation

Aggregate RTCP provides the overall media statistics and follows the standard RTCP defined in RFC3550 [1]. However, subflow specific RTCP provides the per path media statistics because the aggregate RTCP report may not provide sufficient per path information to an MPRTTP scheduler. Specifically, the scheduler should be aware of each path's RTT and loss-rate, which an aggregate RTCP cannot provide. The sender/receiver MUST use non-compound RTCP reports defined in RFC5506 [5] to transmit the aggregate and subflow-specific RTCP reports. Also, each subflow and the aggregate RTCP report MUST follow the timing rules defined in [4].

The RTCP reporting interval is locally implemented and the scheduling of the RTCP reports may depend on the the behavior of each path. For instance, the RTCP interval may be different for a passive path than an active path to keep port bindings alive. Additionally, an endpoint may decide to share the RTCP reporting bit rate equally across all its paths or schedule based on the receiver rate on each path.

8.5. RTCP Transmission

The sender sends an RTCP SR on each active path. For each SR the receiver gets, it echoes one back to the same IP address-port pair that sent the SR. The receiver tries to choose the symmetric path and if the routing is symmetric then the per-path RTT calculations will work out correctly. However, even if the paths are not symmetric, the sender would at maximum, under-estimate the RTT of the path by a factor of half of the actual path RTT.

9. Packet Formats

In this section we define the protocol structures described in the previous sections.

9.1. RTP Header Extension for MPRTTP

The MPRTTP header extension is used to 1) distribute a single RTP stream over multiple subflows, 2) perform connectivity checks on the advertised interfaces, and 3) keep-alive passive interfaces (paths).

The header conforms to the one-byte RTP header extension defined in [6]. The header extension contains a 16-bit length field that counts the number of 32-bit words in the extension, excluding the four-octet extension header (therefore zero is a valid length, see Section 5.3.1 of [1] for details).

To signal the use of the above RTP header extensions in SDP, the following URI MUST be used: urn:ietf:params:rtp-hdext:mp RTP.

The RTP header for each subflow is defined below:

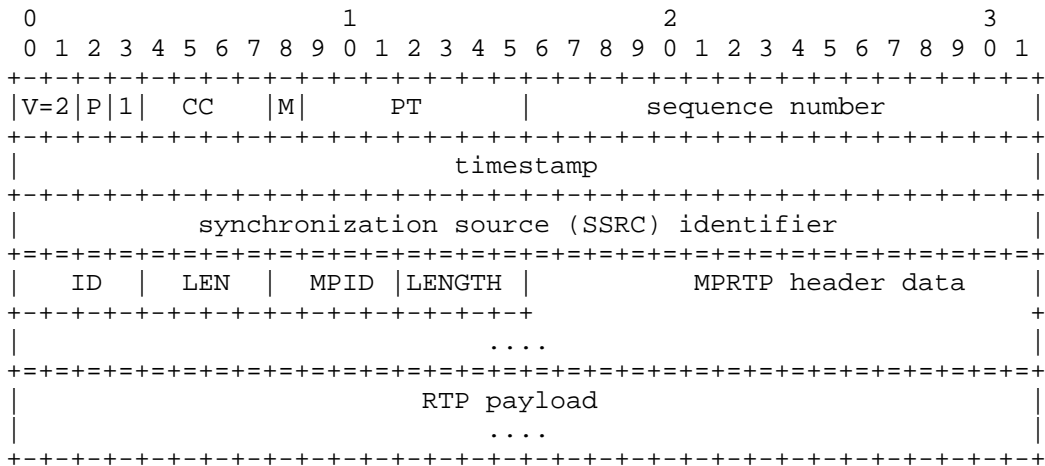


Figure 6: Generic MP RTP header extension

MPID:

The MPID field corresponds to the type of MP RTP packet. Namely:

MPID ID Value	Use
0x00	Subflow RTP Header. For this case the Length is set to 6
0x01	Connectivity Check. For this case the length is set to 0
TBD	Keep Alive Packet.

Figure 7: RTP header extension values for MP RTP (H-Ext ID)

length

The 4-bit length field is the length of extension data in bytes not including the H-Ext ID and length fields. The value zero indicates there is no data following.

MPRTP header data

Carries the MPID specific data as described in the following sub-sections.

9.1.1.1. MPRTP RTP Extension for a Subflow

The RTP header for each subflow is defined below:

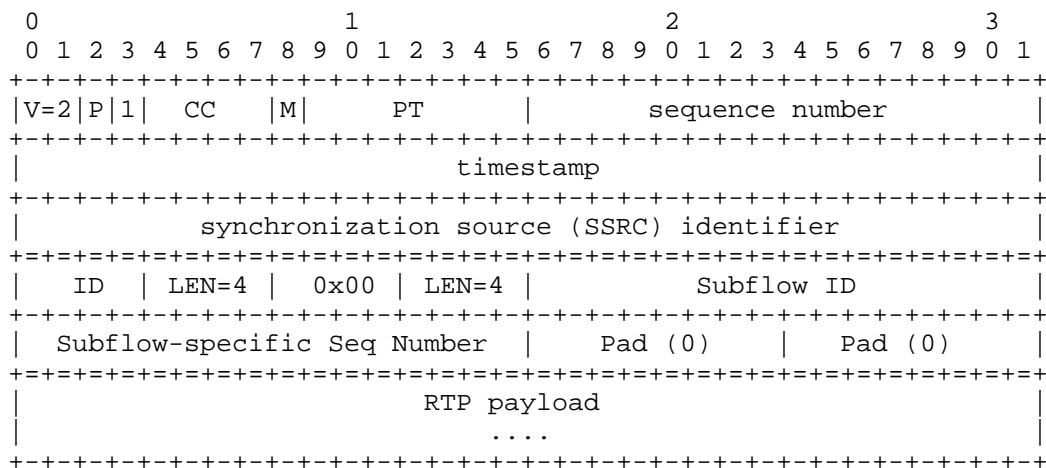


Figure 8: MPRTP header for subflow

MP ID = 0x00

Indicates that the MPRTP header extension carries subflow specific information.

length = 4

Subflow ID: Identifier of the subflow. Every RTP packet belonging to the same subflow carries the same unique subflow identifier.

Flow-Specific Sequence Number (FSSN): Sequence of the packet in the subflow. Each subflow has its own strictly monotonically increasing sequence number space.

9.1.2. MPRTP RTP Extension for Connectivity Checks

[Open Issue: What sequence number to use for the RTP session?
Alternative 1: An MPRTP receiver MUST NOT give the packet with MPID=0x01 to the decoder and ignore these packets from RTCP calculation. Alternative 2: Instead of sending an RTP packet the sender transmits a modified STUN packet.]

9.1.3. MPRTP RTP Extension for Keep-alive Packets

[Editor: RTCP guidelines for keep alive packet [17] recommends multiplexing RTP and RTCP. If so, we can do the same and remove the keep alive from the list.]

9.2. RTCP Extension for MPRTP (MPRTCP)

The MPRTP RTCP header extension is used to 1) provide RTCP feedback per subflow to determine the characteristics of each path, 2) advertise each interface and perform connectivity check on the other endpoint's interfaces, and 3) to keep alive a passive connection.

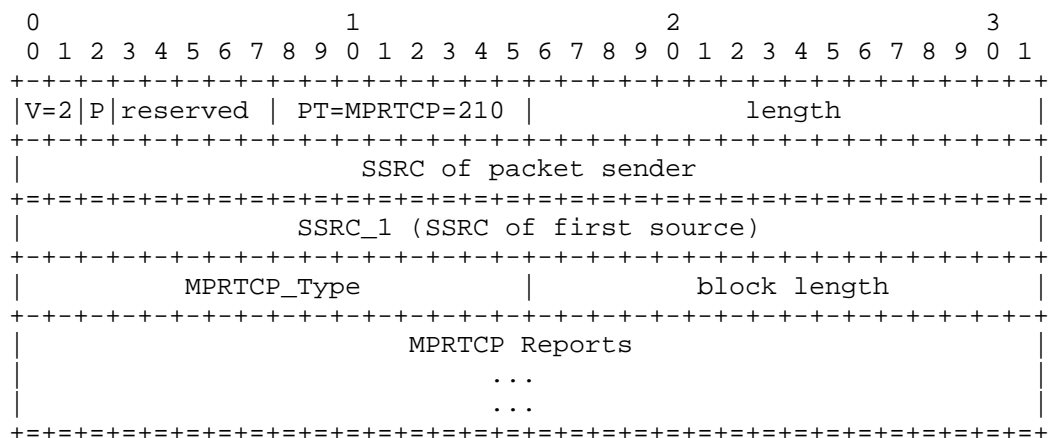


Figure 9: Generic RTCP Extension for MPRTP (MPRTCP) [appended to standard SR/RR]

MPRTCP: 8 bits

Contains the constant 210 to identify this as an Multipath RTCP packet.

length: 16 bits

As described for the RTCP packet (see Section 6.4.1 of the RTP specification [1]), the length of this is in 32-bit words minus one, including the header and any padding.

MPRTCP_Type: 16-bits

The MPRTCP_Type field corresponds to the type of MPRTCP RTCP packet. Namely:

MPRTCP_Type Value	Use
0x00	Subflow Specific Report
0x01	Connectivity Check. For this case the length is set to 0
0x02	Interface Advertisement
TBD	Keep Alive Packet.

Figure 10: RTP header extension values for MPRTCP (MPRTCP_Type)

block length: 16-bits

The 16-bit length field is the length of the encapsulated MPRTCP reports in 32-bit word length not including the MPRTCP_Type and length fields. The value zero indicates there is no data following.

MPRTCP Reports: variable size

Defined later in 9.2.1 and 9.3.1.

9.2.1. MPRTCP Extension for Subflow Reporting

When sending a report for a specific subflow the sender or receiver MUST add only the reports associated with that 5-tuple. Each subflow is reported independently using the following MPRTCP Feedback header.

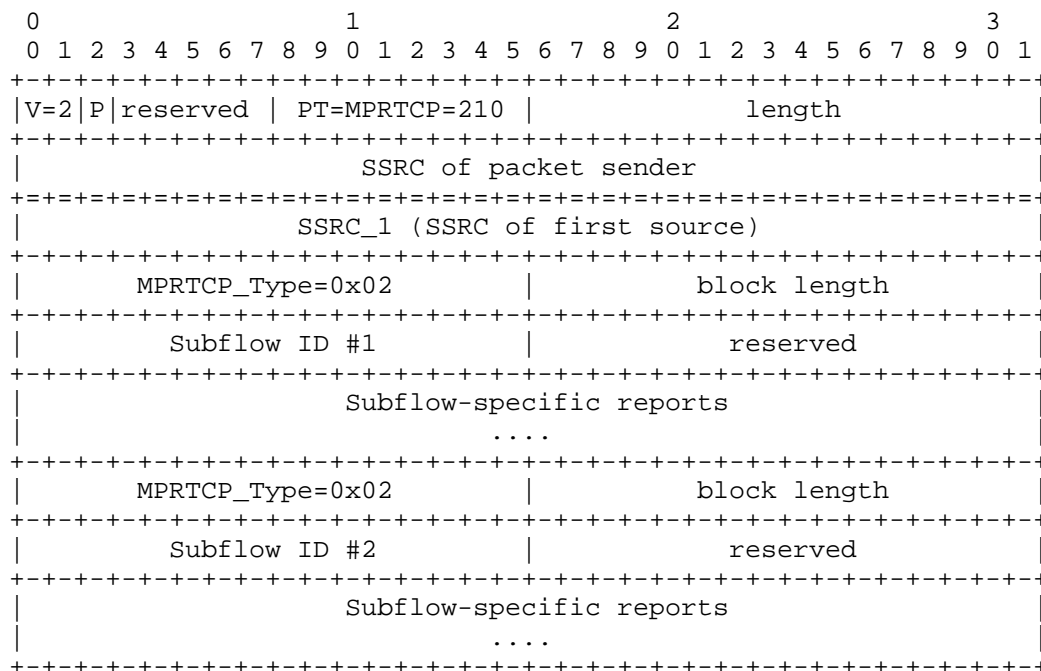


Figure 11: MPRTCP Subflow Reporting Header

MPRTCP_Type: 0x02

block length: 16-bits

The 16-bit length field is the length of the encapsulated subflow-specific reports in 32-bit word length not including the MPRTCP_Type and length fields.

Subflow ID: 16 bits

Subflow identifier is the value associated with the subflow the endpoint is reporting about. If it is a sender it MUST use the Subflow ID associated with the 5-tuple. If it is a receiver it MUST use the Subflow ID received in the Subflow-specific Sender Report.

Subflow-specific reports: variable

Subflow-specific report contains all the reports associated with the Subflow ID. For a sender, it MUST include the Subflow-specific Sender Report (SSR). For a receiver, it MUST include Subflow-specific Receiver Report (SRR). Additionally, if the

receiver supports subflow-specific extension reports then it MUST append them to the SRR.

9.2.1.1. MPRTCP for Subflow-specific SR, RR and XR

[Editor: inside the context of subflow specific reports can we reuse the payload type code for Sender Report (PT=200), Receiver Report (PT=201), Extension Report (PT=207). Transport and Payload specific RTCP messages are session specific and SHOULD be used as before.]

Example:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|V=2|P|reserved | PT=MPRTCP=210 |                length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                SSRC of packet sender |
+=====+=====+=====+=====+=====+=====+=====+=====+
|                SSRC_1 (SSRC of first source) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      MPRTCP_Type=0x02 |                block length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Subflow ID #1 |                reserved |
+=====+=====+=====+=====+=====+=====+=====+=====+
|V=2|P|      RC |      PT=SR=200 |                length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                SSRC of sender |
+=====+=====+=====+=====+=====+=====+=====+=====+
|                NTP timestamp, most significant word |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                NTP timestamp, least significant word |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                RTP timestamp |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                subflow's packet count |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                subflow's octet count |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      MPRTCP_Type=0x02 |                block length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      Subflow ID #2 |                reserved |
+=====+=====+=====+=====+=====+=====+=====+=====+
|V=2|P|      RC |      PT=RR=201 |                length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                SSRC of packet sender |
+=====+=====+=====+=====+=====+=====+=====+=====+
| fraction lost |      cumulative number of packets lost |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                extended highest sequence number received |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                inter-arrival jitter |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                last SR (LSR) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                delay since last SR (DLSR) |
+=====+=====+=====+=====+=====+=====+=====+=====+
|                Subflow specific extension reports |
|                .... |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

Figure 12: Example of reusing RTCP SR and RR inside an MPRTCP header (Bi-directional use-case, in case of uni-directional flow the subflow will only send an SR or RR).

9.3. MPRTCP Extension for Interface advertisement

This sub-section defines the RTCP header extension for in-band interface advertisement by the receiver, instead of relying on ICE or in situations when the interface appears after SDP session establishment.

The interface advertisement SHOULD immediately follow the Receiver Report. If the Receiver Report is not present, then it MUST be appended to the Sender Report.

The endpoint MUST advertise all its interfaces when a new interface appears. Furthermore, an endpoint MUST advertise all its interfaces when it receives an Interface Advertisement.

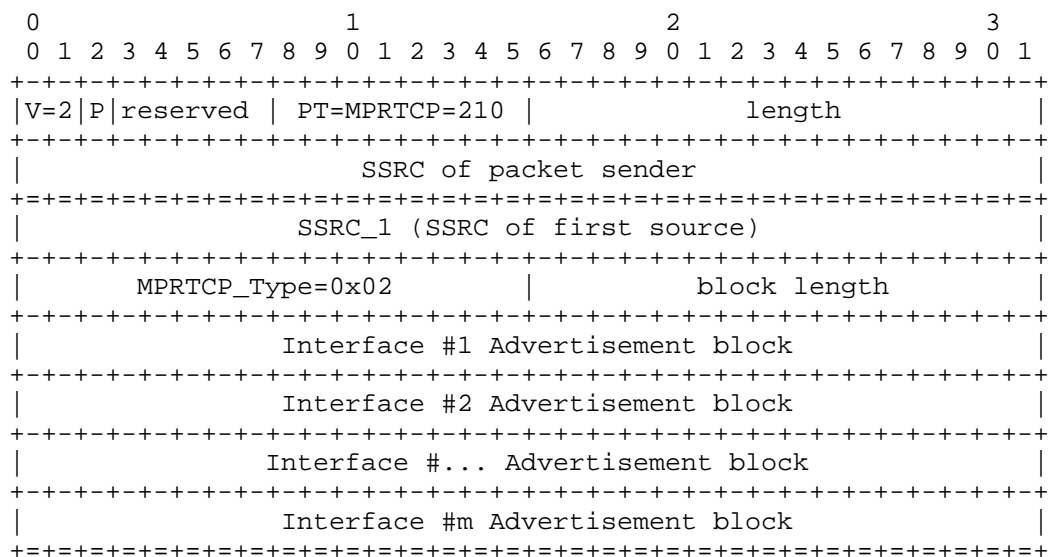


Figure 13: MPRTCP Interface Advertisement. (appended to SR/RR)

MPRTCP_Type: 0x00

block length: 16-bits

The 16-bit length field is the length of the encapsulated interface advertisement blocks in 32-bit word length not including the MPRTCP_Type and length fields.

Interface Advertisement block: variable size
Defined later in 9.3.1.

9.3.1. Interface Advertisement block

This block describes a method to represent IPv4, IPv6 and generic DNS-type addresses in a block format. It is based on the sub-reporting block in [7].




Figure 14: Interface Advertisement block during path discovery

Type: 8 bits

The Type corresponds to the type of address. Namely:

- 0: IPv4 address
- 1: IPv6 address
- 2: DNS name

Length: 8 bits

The length of the Interface Advertisement block in bytes.

For an IPv4 address, this should be 9 (i.e., 5 octets for the header and 4 octets for IPv4 address).

For an IPv6 address, this should be 21.

For a DNS name, the length field indicates the number of octets making up the string plus the 5 byte header.

RTP Port: 2 octets

The port number to which the sender sends RTP data. A port number of 0 is invalid and MUST NOT be used.

RTCP Port: 2 octets

The port number to which receivers send feedback reports. A port number of 0 is invalid and MUST NOT be used.

Address: 4 octets (IPv4), 16 octets (IPv6), or n octets (DNS name)

The address to which receivers send feedback reports. For IPv4 and IPv6, fixed-length address fields are used. A DNS name is an arbitrary-length string. The string MAY contain Internationalizing Domain Names in Applications (IDNA) domain names and MUST be UTF-8 [8] encoded.

10. RTCP Timing reconsiderations for MPRTCP

MPRTCP endpoints MUST conform to the timing rule imposed in [4], i.e., the total RTCP rate between the participants MUST NOT exceed 5% of the media rate. For each endpoint, a subflow MUST send the aggregate and subflow-specific report. The endpoint SHOULD schedule the RTCP reports for the active subflows based on the share of the transmitted or received bit rate to the average media bit rate, this method ensures fair sharing of the RTCP bandwidth. Alternatively, the endpoint MAY schedule the reports in round-robin.

11. SDP Considerations

The packet formats specified in this document define extensions for RTP and RTCP. The use of MPRTCP is left to the discretion of the sender and receiver.

11.1. Signaling MPRTCP capability in SDP

A participant of a media session MAY use SDP to signal that it supports MPRTCP. Not providing this information may/will make the sender or receiver ignore the header extensions. However, MPRTCP MAY be used by either sender or receiver without prior signaling.

```
mprtcp-attrib = "a=" "mprtcp" [  
    SP mprtcp-optional-parameter]  
    CRLF ; flag to enable MPRTCP
```

The literal 'mprtp' MUST be used to indicate support for MPRTTP. Generally, senders and receivers MUST indicate this capability if they support MPRTTP and would like to use it in the specific media session being signaled. They can then use the in-band signaling to exchange the additional interfaces and to perform connectivity checks.

11.2. Signaling MPRTTP Header Extension in SDP

To signal the use of MPRTTP header extensions in SDP, the following URI MUST be used: urn:ietf:params:rtp-hdext:mprtp. Legacy RTP (non-MPRTTP) clients MUST ignore this header extension, but continue to parse and decode the packet.

11.3. MPRTTP using preloaded interfaces from ICE (out-band signaling)

If the endpoints intend to use ICE [3] for discovering new interfaces and running connectivity checks then these endpoints MUST NOT use in-band interface advertisement and the associated connectivity checks, as described in Section 9.1.2 (MPID=0x01) and Section 9.3 (MPRTCP_Type=0x01, 0x02).

Using ICE for MPRTTP is a two step process:

1. Advertising ICE candidates: in the initial INVITE the endpoints exchange candidates, as defined in ICE [3]. Thereafter the endpoints run connectivity checks
2. Advertising MPRTTP interfaces: When a sufficient number of connectivity checks succeed the endpoint MUST send an updated offer containing their interfaces.

This process enables the participants to use MPRTTP capabilities from the start of a media session.

11.3.1. "interface" attribute

The interface attribute is only a media-level attribute. It is used to describe one of the many network addresses of an endpoint.

The syntax of the interface attribute is defined using the following Augmented BNF, as defined in [9]. It reuses the RFC4566 [18] definitions of unicast-address, port, token, SP, and CRLF.

```
mprtp-optional-parameter = mprtp-interface
                           ; other optional parameters may be added later

mprtp-interface = "interface" ":" counter SP unicast-address
                  ":" rtp_port "/" rtcp_port
                  *(SP interface-description-extension)

counter = 1*DIGIT
rtp_port = port      ;port from RFC4566
rtcp_port = port     ;port from RFC4566
```

<mprtp-interface>: specifies one of the unicast IP address, the RTP and RTCP port numbers of the endpoint. The unicast address in the first interface line MUST match the connection address ('c=' line). The RTP and RTCP ports in the first interface line MUST match with the RTP and RTCP ports in the associated 'm=' line. The counter should start at 1 and increment with each additional interface. Multiple interface lines MUST be ordered in a decreasing priority level as is the case with the Interface Advertisement blocks in in-band signaling (See Figure 13).

<unicast-address>: is taken from RFC5245 [3]. It is one of the IP addresses of the endpoint allows the use of IPv4 addresses, IPv6 addresses and Fully Qualified Domain Names (FQDN). An endpoint MUST only include the IP address for which the connectivity checks have succeeded.

<port>: is from RFC4566 [18]. It is the RTP or RTCP port associated the unicast address for which the connectivity checks have succeeded.

<counter>: is a monotonically increasing positive integer starting at 1. The counter SHOULD reset for each media line.

The 'mprtp-interface' can be extended using the 'interface-description-extension' parameter. An endpoint MUST ignore any extensions it does not understand.

11.3.2. Example

The ABNF grammar is illustrated by means of an example:

```
v=0
o=alice 2890844526 2890844527 IN IP4 192.0.2.1
s=
c=IN IP4 192.0.2.1
t=0 0
m=video 49170 RTP/AVP 98
a=rtpmap:98 H264/90000
a=fmtp:98 profile-level-id=42A01E;
a=mp RTP interface:1 192.0.2.1:49170/49171 ;primary interface
a=mp RTP interface:2 192.1.2.1:51372/51373 ;additional interface
```

11.4. Increased Throughput

The MP RTP layer MAY choose to augment paths to increase throughput. If the desired media rate exceeds the current media rate, the endpoints MUST renegotiate the application specific ("b=AS:") [18] bandwidth.

11.5. Offer/Answer Examples

This section shows examples of SDP offer and answer for in-band and out-band signaling.

11.5.1. In-band signaling

The following offer/answer shows that both the endpoints are MP RTP capable and SHOULD use in-band signaling for exchanging interfaces and for performing connectivity checks. If candidate addresses are included in the initial INVITE then the connectivity checks are performed by ICE and the out-band signaling is preferred.

```
Offer:
v=0
o=alice 2890844526 2890844527 IN IP4 192.0.2.1
s=
c=IN IP4 192.0.2.1
t=0 0
m=video 49170 RTP/AVP 98
a=rtpmap:98 H264/90000
a=fmtp:98 profile-level-id=42A01E;
a=mp RTP
```

Answer:

```
v=0
o=bob 2890844528 2890844529 IN IP4 192.0.2.2
s=
c=IN IP4 192.0.2.2
t=0 0
m=video 4000 RTP/AVP 98
a=rtpmap:98 H264/90000
a=fmtp:98 profile-level-id=42A01E;
a=mp RTP
```

11.5.2. Out-band signaling

In this example, the offerer advertises two interfaces and the answerer responds with a single interface description. The endpoint MAY use one or both paths depending on the end-to-end characteristics of each path.

Offer:

```
v=0
o=alice 2890844526 2890844527 IN IP4 192.0.2.1
s=
c=IN IP4 192.0.2.1
t=0 0
m=video 49170 RTP/AVP 98
a=rtpmap:98 H264/90000
a=fmtp:98 profile-level-id=42A01E;
a=mp RTP interface:1 192.0.2.1:49170/49171
a=mp RTP interface:2 192.1.2.1:51372/51373
```

Answer:

```
v=0
o=bob 2890844528 2890844529 IN IP4 192.0.2.2
s=
c=IN IP4 192.0.2.2
t=0 0
m=video 4000 RTP/AVP 98
a=rtpmap:98 H264/90000
a=fmtp:98 profile-level-id=42A01E;
a=mp RTP interface:1 192.0.2.2:4000/4001
```

12. IANA Considerations

TBD

13. Security Considerations

TBD

All drafts are required to have a security considerations section. See RFC 3552 [19] for a guide.

14. Acknowledgements

Varun Singh, Saba Ahsan, and Teemu Karkkainen are supported by Trilogy (<http://www.trilogy-project.org>), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Program. The views expressed here are those of the author(s) only. The European Commission is not liable for any use that may be made of the information in this document.

The authors would also like acknowledge the contribution of Ralf Globish and Thomas Schierl for providing the input and text for the MPRTTP interface advertisement in SDP.

15. References

15.1. Normative References

- [1] Schulzrinne, H., Casner, S., Frederick, R., and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", STD 64, RFC 3550, July 2003.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [3] Rosenberg, J., "Interactive Connectivity Establishment (ICE): A Protocol for Network Address Translator (NAT) Traversal for Offer/Answer Protocols", RFC 5245, April 2010.
- [4] Ott, J., Wenger, S., Sato, N., Burmeister, C., and J. Rey, "Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)", RFC 4585, July 2006.
- [5] Johansson, I. and M. Westerlund, "Support for Reduced-Size Real-Time Transport Control Protocol (RTCP): Opportunities and Consequences", RFC 5506, April 2009.
- [6] Singer, D. and H. Desineni, "A General Mechanism for RTP Header Extensions", RFC 5285, July 2008.

- [7] Ott, J., Chesterfield, J., and E. Schooler, "RTP Control Protocol (RTCP) Extensions for Single-Source Multicast Sessions with Unicast Feedback", RFC 5760, February 2010.
- [8] Yergeau, F., "UTF-8, a transformation format of ISO 10646", STD 63, RFC 3629, November 2003.
- [9] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.

15.2. Informative References

- [10] Stewart, R., "Stream Control Transmission Protocol", RFC 4960, September 2007.
- [11] Ford, A., Raiciu, C., Handley, M., Barre, S., and J. Iyengar, "Architectural Guidelines for Multipath TCP Development", RFC 6182, March 2011.
- [12] Nordmark, E. and M. Bagnulo, "Shim6: Level 3 Multihoming Shim Protocol for IPv6", RFC 5533, June 2009.
- [13] Westerlund, M. and S. Wenger, "RTP Topologies", RFC 5117, January 2008.
- [14] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [15] Schulzrinne, H., Rao, A., Lanphier, R., Westerlund, M., and M. Stiemerling, "Real Time Streaming Protocol 2.0 (RTSP)", draft-ietf-mmusic-rfc2326bis-27 (work in progress), March 2011.
- [16] Rosenberg, J. and H. Schulzrinne, "An Offer/Answer Model with Session Description Protocol (SDP)", RFC 3264, June 2002.
- [17] Marjou, X. and A. Sollaud, "Application Mechanism for Keeping Alive the NAT Mappings Associated with RTP / RTP Control Protocol (RTCP) Flows", RFC 6263, June 2011.
- [18] Handley, M., Jacobson, V., and C. Perkins, "SDP: Session Description Protocol", RFC 4566, July 2006.
- [19] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, July 2003.

Authors' Addresses

Varun Singh
Aalto University
School of Science and Technology
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: varun@comnet.tkk.fi
URI: <http://www.netlab.tkk.fi/~varun/>

Teemu Karkkainen
Aalto University
School of Science and Technology
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: teemuk@comnet.tkk.fi

Joerg Ott
Aalto University
School of Science and Technology
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: jo@comnet.tkk.fi

Saba Ahsan
Aalto University
School of Science and Technology
Otakaari 5 A
Espoo, FIN 02150
Finland

Email: sahsan@cc.hut.fi

Lars Eggert
Nokia Research Center
P.O. Box 407
Nokia Group 00045
Finland

Phone: +358 50 48 24461
Email: lars.eggert@nokia.com
URI: http://research.nokia.com/people/lars_eggert

