

RTP with TCP Friendly Rate Control

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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>.

Copyright Notice

Copyright (C) The Internet Society (2007).

Abstract

This memo specifies how the TCP Friendly Rate Control (TFRC) of RTP flows can be supported using the RTP/AVPF profile and the general RTP header extension mechanism. AVPF feedback packets and RTP header extensions are defined to support the exchange of control information between RTP TFRC senders and receivers. TFRC is an equation based congestion control scheme for unicast flows operating in a best effort Internet environment.

1. Introduction

[Note to RFC Editor: All references to RFC XXXX are to be replaced with the RFC number of this memo, when published]

This memo specifies how the TCP Friendly Rate Control (TFRC) of RTP flows can be supported using the RTP/AVPF profile and RTP header extensions, by defining the header extensions to be used and a new AVPF feedback packet.

TFRC is an equation based congestion control scheme for unicast flows operating in a best effort Internet environment and competing with TCP traffic. TFRC computes a TCP-friendly data rate based on current network conditions, as represented by the latest round trip time and packet loss calculations. The complete TFRC mechanism is described in detail in [[TFRC](#)].

To calculate a TCP-friendly data rate and keep track of round trip times and packet losses, TFRC senders and receivers rely on exchanging specific information between each other, i.e: the sender provides the receiver with the latest updates to round trip time calculations, while the receiver provides feedback needed to compute round trip times and on packet losses. This memo defines how this information can be exchanged between TFRC senders and receiver with RTP header extensions and an AVPF feedback packet.

2. Terminology

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 [RFC 2119](#) [[2119](#)].

3. Relation to the Datagram Congestion Control Protocol

The TFRC congestion control mechanism is also supported by the Datagram Congestion Control Protocol (DCCP). In this section we detail the pros and cons of using TFRC with RTP versus DCCP.

DCCP is a minimal general purpose transport-layer protocol with unreliable yet congestion controlled packet delivery semantics and reliable connection setup and teardown. DCCP currently supports both TFRC and TCP-like congestion control, and the protocol is structured to support new congestion control mechanisms defined in the future. In addition DCCP supports a host of other features, such as: use of Explicit Congestion Notification (ECN) and the ECN Nonce, reliable option negotiation and Path Maximum Transfer Unit (PMTU). Naturally an application using RTP/DCCP as its transport protocol will benefit from the protocol features supported by DCCP.

However there are a number of benefits to be gained by the development and standardization of the use RTP with TFRC:

- o Media applications lacking congestion control can incorporate congestion controlled transport without delay by using RTP with TFRC. The DCCP protocol is currently under development and widespread deployment is not yet in place.
- o Use of the RTP with TFRC is not contingent on any OS level changes and can be quickly deployed, as RTP is implemented at the application layer.
- o RTP/UDP flows face the same restrictions in firewall traversal as do UDP flows and do not require NATs and firewall modifications. DCCP flows, on the other hand, do require NAT and firewall modifications, however once these modifications are in place, they can result in easier NAT and firewall traversal for RTP/DCCP flows in the future.
- o Use of RTP with TFRC with various media applications will give researchers, implementors and developers a better understanding of the intricate relationship between media quality and equation based congestion control. Hopefully this experience with congestion control and TFRC will ease the migration of media applications to DCCP once DCCP is deployed.

Overall, using the AVPF/RTP profile and header extension to support TFRC provides an immediate means for congestion control in media streams, in the time being until DCCP is deployed.

Additionally, there are also a number of technical differences as to how (and which) congestion control information is exchanged between DCCP with CCID3 and RTP:

- o Using header extensions the RTP TFRC sender transmits a send timestamp to the RTP TFRC receiver with every data packet. In addition to congestion control the send timestamp can be used by the receiver for jitter calculations.

In contrast DCCP with CCID3 transmits a quad round trip counter to the receiver.

- o The RTP TFRC receiver only provides the RTP TFRC sender with the loss event rate as computed by the receiver.

In contrast DCCP with CCID3, provides 2 other options for the transport of loss event rate. A sender may choose to receive loss intervals or an Ack Vector. These two options provide the

sender with the necessary information to compute the loss event rate.

- o Sequence number: DCCP supports a 48 bit and a 24 bit sequence number, whereas RTP only supports a 16 bit sequence number. While this makes RTP susceptible to data injection attacks, it can be avoided by using the SRTP [[SRTP](#)] profile.

4. The TFRC Information Exchange Loop

TFRC depends on the exchange of congestion control information between a sender and receiver. In this section we reiterate which items are exchanged between a TFRC sender and receiver as discussed in [[TFRC](#)]. We note how the RTP can accommodate these exchanges.

4.1. Data Packets

As stated in [[TFRC](#)] a TFRC sender transmits the following information in each data packet to the receiver:

- o A sequence number, incremented by one for each data packet transmitted.
- o A timestamp indicating the packet send time and the sender's current estimate of the round-trip time, RTT. This information is then used by the receiver to compute the TFRC loss intervals.
 - or -
 - A course-grained timestamp incrementing every quarter of a round trip time, which is then used to determine the TFRC loss intervals.

The standard RTP sequence number suffices for TFRCs functionality. RTP header extension [hdrtxt] are used to transmit the send timestamp and RTT. The RTT can be transmitted in band with every RTP packet or when there is significant change in the RTT. Each extension payload is 3 bytes long (see [Section 6](#)).

4.2. Feedback Packets

As stated in [[TFRC](#)] a TFRC receiver provides the following feedback to the sender at least once per RTT or per data packet received (whichever time interval is larger):

- o The send timestamp of the last data packet received, t_i .
- o The amount of time elapsed between the receipt of the last data packet at the receiver, and the generation of this feedback

report, `t_delay`. This is used by the sender for RTT computations.

- o The rate at which the receiver estimates that data was received since the last feedback report was sent, `x_recv`.
- o The receiver's current estimate of the loss event rate, `p`, a real value between 0 and 1.0.

To accommodate the feedback of these values a new AVPF transport layer feedback message is defined, as detailed in [Section 7](#).

5. The Header Extensions

The form of the extension block when both the RTT and send timestamp are being transmitted is depicted in Figure 1. The length field for each extensions takes the value 2 to indicate that the payload is 3 bytes. The two header extension fields are defined and used as follows:

Send timestamp: 24 bits

The timestamp indicating when the packet is sent. This timestamp is measured in microseconds and is used for round trip time calculations.

Round trip time (RTT): 24 bits

The round trip time as measured by the RTP TFRC sender in microseconds.



Figure 1

6. TFRC-FB: A New AVPF Transport Layer Feedback Message

To support feedback to the receivers a new transport layer AVPF feedback message is defined: TFRC-FB. This message is depicted in Figure 2. It is defined according to [\[AVPF\]](#) and includes the following four fields:

Timestamp (t_i): 32 bits

The send timestamp of the last data packet received by the RTP TFRC receiver, t_i , in microseconds.

Delay (t_{delay}): 32 bits

The amount of time elapsed between the receipt of the last data packet at the RTP TFRC receiver, and the generation of this feedback report in microseconds. This is used by the RTP TFRC sender for RTT computations.

Data rate (x_{recv}): 32 bits

The rate at which the receiver estimates that data was received since the last feedback report was sent in bytes per second.

Loss event rate (p): 32 bits

The receiver's current estimate of the loss event rate, p , expressed as a fixed point number with the binary point at the left edge of the field. (That is equivalent to taking the integer part after multiplying the loss event rate by 2^{32} .)

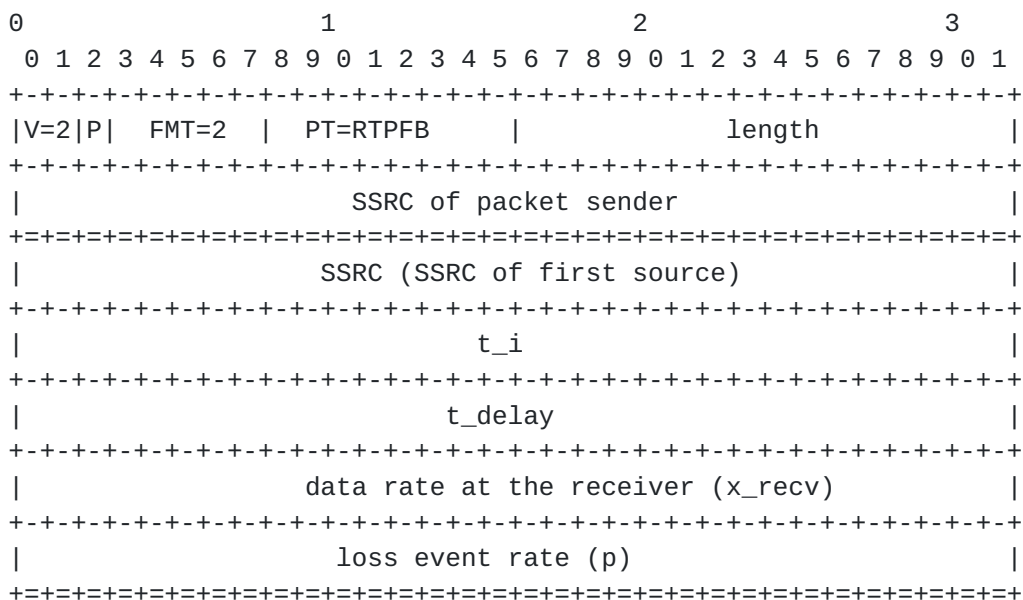


Figure 2

7. RTCP Transmission Intervals, TFRC and the AVPF Profile

When running TFRC rate controlled RTP, the RTCP transmission intervals MUST be set according to the requirements of the TFRC algorithm. TFRC requires a receiver to generate a feedback packet at least once per RTT or per packet received (based on the larger time

interval). These requirements are to ensure timely reaction to congestion.

To support the transmission of feedback packet once per RTT, a RTP/AVPF flow with TFRC congestion control must:

- o set `allow_early` to "true" at all times. Essentially, this means that a receiver can always generate an early feedback packet, and does not need to alternate between early and regular RTCP packets (see [RFC 4585, Section 3.4,k](#)).
- o `T_rr_interval` must not be set to a value larger than the current round trip time, as this would prevent generating feedback packets at least once per RTT (see [RFC 4585, Section 3.4,m](#)).

The TFRC requirements of receiving feedback once per RTT can at times conflict with the AVP RTCP bandwidth constraints, particularly at small RTTs of 20ms or less. Assuming only one TFRC-FB report per RTCP compound packet, Table 1 lists the RTCP bandwidths at RTTs of 2, 5, 10 and 20 ms and the minimum corresponding RTP data rates, where $RTCP(X) \leq (0.05) * RTP(X)$ is true. For example, according to Table 1, a TFRC RTP flow of less than 3.2 Mbps and a RTT of 5 ms, can not comply with the 5% RTCP bandwidth constraints (Table 1 assumes each RTCP packet is 100 bytes). RTP flows facing such circumstance should take into account the additional RTCP bandwidth needed when signaling their bandwidth information in SDP.

RTT	RTCP(X)	RTP(X)
+-----+		
20 ms	40 kbps	0.8 Mbps
10 ms	80 kbps	1.6 Mbps
5 ms	160 kbps	3.2 Mbps
2 ms	400 kbps	8.0 Mbps
+-----+		

Table 1

8. SDP Definitions

RTP flows using TFRC congestion control must signal their use of header extensions for round trip times (RTT) and the send timestamp:

```
a=extmap:4 urn:ietf:params:rtp-hdtext:rtt
a=extmap:4 urn:ietf:params:rtp-hdtext:send-ts
```


9. IANA Considerations

In this section we detail IANA registry values that need to be registered.

The new RTP/AVPF feedback packet, TFRC-FB, must be registered. For the RTPFB range of packets, the following format (FMT) values are needed:

Value name: TFRC-FB
Long name: TFRC feedback
Value: 2
Reference: RFC XXXX

The names rtt and send-ts need to be registered into the rtp-hdrex section of the urn:ietf: namespace, referring to RFC XXXX.

10. Security Considerations

This memo defines how to use the RTP AVPF profile and the general RTP header extensions to support TFRC congestion control. Therefore RTP packets using these mechanisms are subject to the security considerations discussed in the RTP specification [[RTP](#)], the RTP/AVPF profile specification [[AVPF](#)] and the general header extensions mechanism [hdrtxt]. Combining these mechanisms does not pose any additional security implications. Applications requiring authentication and integrity protection can use the SAVPF [[SAVPF](#)] profile.

11. Acknowledgments

This memo is based upon work supported by the U.S. National Science Foundation (NSF) under Grant No. 0334182. Any opinions, findings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of NSF.

12. Author's Address

Ladan Gharai <ladan@isi.edu>
USC Information Sciences Institute
3811 N. Fairfax Drive, #200
Arlington, VA 22203
USA

Normative References

- [RTP] H. Schulzrinne, S. Casner, R. Frederick and V. Jacobson, "RTP: A Transport Protocol for Real-Time Applications", Internet Engineering Task Force, [RFC 3550](#) (STD0064), July 2003.
- [AVP] H. Schulzrinne and S. Casner, "RTP Profile for Audio and Video Conferences with Minimal Control," [RFC 3551](#) (STD0065), July 2003.
- [AVPF] J. Ott, S. Wenger, A. Sato, C. Burmeister and J. Ray, "Extended RTP Profile for RTCP-based Feedback (RTP/AVPF)", [RFC 4585](#), July 2006.
- [2119] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", Internet Engineering Task Force, [RFC 2119](#), March 1997.
- [2434] T. Narten and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", Internet Engineering Task Force, [RFC 2434](#), October 1998.
- [TFRC] M. Handley, S. Floyd, J. Padhye and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", Internet Engineering Task Force, [RFC 3448](#), January 2003.
- [SDP] M. Handley and V. Jacobson, "SDP: Session Description Protocol", [RFC 2327](#), April 1998.
- [SRTP] M. Baugher, D. McGrew, M. Naslund, E. Carrara, K. Norrman, "The Secure Real-time Transport Protocol", [RFC 3711](#), March 2004.
- [hdrext] D. Singer, "A general mechanism for RTP Header Extensions", ID [draft-ietf-avt-rtp-hdrext-08](#), October 2006.
- [SAVPF] J. Ott, E. Carrara, "Extended Secure RTP Profile for RTCP-based Feedback (RTP/SAVPF)," [draft-ietf-avt-profile-savpf-09.txt](#), April, 2007.

Informative References

[13. IPR Notice](#)

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to

pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in [BCP 78](#) and [BCP 79](#).

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

14. Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in [BCP 78](#), and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."

