Audio/Video Transport Working Group
Internet Draft
Document:draft-fukunaga-low-delay-rtcp-02.txt
February 01, 2001
Expires: August 01, 2001

Shigeru Fukunaga
Noriyuki Sato
Oki

Koichi Yano
FastForward Networks

Akihiro Miyazaki
Koichi Hata
Rolf Hakenberg
Carsten Burmeister
Matsushita

Low Delay RTCP Feedback Format

Status of this Memo

1. Abstract

Low delay feedback in RTP has gained a lot of interest recently in order to enhance the performance in RTP sessions with two or only a small number of participants. A recent document describes rules that enhance the functionality of the existing RTP timing and allow low delay feedback. While that document describes when it is allowed to send data, we describe here what kind of low delay feedback may be sent. Therefore we define three general RTCP feedback message types as well as some recommendations how the sender SHOULD react to certain feedback. The message types can either be used as is or they can be extended in payload or application specific ways.

2. Conventions used in this document

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 [2].

3. Introduction

RTP and its profile define strict rules on the timing behavior of RTCP feedback. It is very important for large multicast groups to define such rules, because the traffic on the feedback channel might explode otherwise. However for small multicast groups or even unicast, a frequently used application of RTP, the rules can be less strict.

A new feedback timing behavior for RTP is defined in [8]. The document describes in certain scenarios when it is allowed to send feedback. These new timing rules scale from unicast connections, where feedback is possible anytime immediately to small multicast groups, where feedback can be sent with low delay according to a credit based scheduling scheme.

While [8] describes the behavior when the receiver is allowed to send feedback, this document describes a framework what is going to be sent. Section 4 describes three general feedback packet types. General formats and semantics are given, but room is left for the extension to specific feedback messages. Section 5 describes some recommendations how the sender SHOULD react if certain feedback is received. Especially congestion control items are discussed.

4 Low Delay Feedback Packets

In this document we define a framework of low delay RTCP feedback messages. Therefore three different types of messages are defined, which are categorized as follows:

- Transport Layer Feedback Messages
- Payload Specific Feedback Messages
- Application Layer Feedback Messages

Transport Layer Feedback Messages are used to transmit protocol related information from the receiver to the sender. This information could be a general acknowledgement (ACK) or negative-acknowledgement (NACK). It could be thought of other types of Transport Layer Feedback Messages, which can be registered later then.

Payload Specific Feedback Messages transport information that is specific to certain payloads. The format and type of information is defined either directly within an RTP Payload Format or in additional feedback format documents.

Application Layer Feedback Messages are used to transport application specific feedback directly from the receiver's application to the sender's application. These data that is sent between the applications is usually defined in the application's specification, e.g. NEWPRED messages are defined in MPEG-4 specification. The data can be identified by the application and therefore needs no additional external information.

The general syntax and semantics for the three packet types is described in the following subsections.

4.1 Transport Layer Feedback Messages

This feedback messages are used to inform the sender's RTP protocol entity about certain events that the receiver detected. These events could be e.g. the reception of a packet or the detection of a packet loss.

The generic format of the packet is as follows:

```
    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|E|  FMT  |    PT=RTPFB   |          length               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                  SSRC of packet sender                        |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                  SSRC of media source                         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   :                              FCI                              :
   :                                                               :
```

The various fields V, P, SSRC and length are defined in the RTP specification [3]. We summarize the meaning of all of the fields below:

version (V): 2 bits
This field identifies the RTP version. The current version is 2.

padding (P): 1 bit
If set, the padding bit indicates that the packet contains additional padding octets at the end which are not part of the control information but are included in the length field.

early packet (E): 1 bit
This bit is set if the packet is sent early. A detailed description can be found in [8].

feedback message type (FMT): 4 bits
This field identifies the type of the feedback. The packet can be used for different purposes. The following types are defined (up to now):

- 0: forbidden
- 1: General NACK
- 2: General ACK
- 3-15: reserved

packet type (PT): 8 bits
This is the RTCP packet type which identifies the packet as being an RTP Feedback Message.

length: 16 bits
The length of this packet in 32-bit words minus one, including the header and any padding. This conforms with the definition of the length field used in RTCP sender and receiver reports [3].

SSRC of packet sender: 16 bits
The synchronization source identifier for the originator of this packet.

SSRC of media source: 16 bits
The synchronization source identifier of the media source that this

feedback is related to.

feedback control information (FCI): variable length
Format and semantic of the FCI varies for the different message types. See sections 5.1.1 and 5.1.2 for the definition of the syntax and the semantic of this field for the general NACK and ACK packets.

4.1.1 General NACK

In this section the general NACK packet is described. As said above, the general NACK packet is of the type Transport Layer Feedback Message. The basic header of section 5.1 applies for this packet, where FMT is set to 1 (General NACK).

The general NACK packet is used to indicate the loss of one or several RTP packets. Therefore the lost packet(s) are identified by the means of a packet identifier and a bitmask.

The Feedback control information (FCI) field has the following syntax:


```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |            PID                |             BLP               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Packet ID (PID): 16 bits
The PID field is used to specify a lost packet. RTP Payload Formats may decide how to identify a packet. Typically RTP sequence number is used for PID as the default format.

bitmask of following lost packets (BLP): 15 bits
The BLP allows for reporting losses of any of the 15 RTP packets immediately following the RTP packet indicated by the PID. The BLP's definition is identical to that given in [RFC2032]. Denoting the BLP's least significant bit as bit 1, and its most significant bit as bit 15, then bit i of the bitmask is set to 1 if the sender has not received RTP packet number PID+i (modulo 2^16) and the receiver decides this packet is lost; bit i is set to 0 otherwise. Note that the sender MUST NOT assume that a receiver has received a

packet because its bitmask was set to 0. For example, the least significant bit of the BLP would be set to 1 if the packet corresponding to the PID and the following packet have been lost. However, the sender cannot infer that packets PID+2 through PID+15 have been received simply because bits 2 through 15 of the BLP are 0; all the sender knows is that the receiver has not reported them as lost at this time.

4.1.2 General ACK

The general ACK packet is also of the type Transport Layer Feedback Message and therefore uses the header as defined in section 5.1 with FMT=2 (general ACK).

The general ACK packet is used to indicate that one or several RTP packets are received correctly. Therefore the received packet(s) are identified by the means of a packet identifier and a bitmask. Acking of a range of consecutive packets is also possible.

The Feedback control information (FCI) field has the following syntax:

```
    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
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |          1st PID              |R|          BLP/PID            |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```


Packet ID (1st PID): 16 bits
This PID field is used to specify a correctly received packet. RTP Payload Formats may decide how to identify a packet. Typically RTP sequence number is used for PID as the default format.

Range of ACKs (R): 1 bit
The R-bit indicates that a range of consecutive packets are received correctly. The 1st PID field specifies the first packet of that range and the next field (BLP/PID) will carry the last PID of the range that is acknowledged.

bitmask of following lost packets / Packet ID (BLP/PID): 15 bits
The semantic of this field changes according to the value of the R-

field. If R=1, this field is used to identify the last packet of the acknowledged packet range, as described above.
If R=0, this field carries a bitmask. The BLP allows for reporting reception of any of the 15 RTP packets immediately following the RTP packet indicated by the PID. The BLP's definition is identical to that given in [RFC2032]. Denoting the BLP's least significant bit as bit 1, and its most significant bit as bit 15, then bit i of the bitmask is set to 1 if the sender has received RTP packet number PID+i (modulo 2^16) and the receiver decides to ACK this packet; bit i is set to 0 otherwise.

4.2 Payload Specific Feedback Messages

Packets of the type Payload Specific Feedback Messages are used to transport information from the receiver to the sender that is directly related to the payload of the forward stream. Messages of this type include e.g. Picture or Slice Loss Indications.

The basic header format is described below.

```
 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|E|  FMT  |    PT=PSFB    |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  SSRC of packet sender                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  SSRC of media source                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|-|   RTP PT    |                                               :
+-+-+-+-+-+-+-+-+                      FCI                      :
:                                                               :
```

The various fields V, P, SSRC and length are defined in the RTP specification [3]. We summarize the meaning of all of the fields below:

version (V): 2 bits
This field identifies the RTP version. The current version is 2.

padding (P): 1 bit

If set, the padding bit indicates that the packet contains additional padding octets at the end which are not part of the control information but are included in the length field.

early packet (E): 1 bit
This bit is set if the packet is sent early. A detailed description can be found in [8].

feedback message type (FMT): 4 bits
This field identifies the type of the feedback. Every RTP Payload Format has its own FMT number space, i.e. a certain FMT number can have different meanings for different payload types. It is therefore expected that either the RTP Payload Format documents itself or additional feedback format documents (related to one or several Payload Formats) define the FMT numbers and their meanings. FMT=0 is forbidden.

packet type (PT): 8 bits
This is the RTCP packet type which identifies the packet as being an Payload Specific Feedback Message.

length: 16 bits
The length of this packet in 32-bit words minus one, including the header and any padding. This conforms with the definition of the length field used in RTCP sender and receiver reports [3].

SSRC of packet sender: 16 bits
The synchronization source identifier for the originator of this packet.

SSRC of media source: 16 bits
The synchronization source identifier of the media source that this feedback is related to.

RTP payload type (RTP PT): 7 bits
This field indicates the payload format to which the feedback is related. It is important to use that field, because the session could be used with several payload formats on the forward direction and the semantics of the packet is different for different payload formats.

feedback control information (FCI): variable length
Format and semantic of the FCI varies for the different message types. The type of FCI that is used is indicated by the combination of FMT and RTP PT. The mapping of FMT numbers to FCIs is expected to

be described in additional payload specific documents or the RTP Payload Formats.

4.2.1 Example

This section describes a short example how the payload specific feedback messages can be defined and used.

As an example we consider a predictive coded video application. If packet loss occurs, it is important that the sender gets this information, because it will otherwise send data, which the receiver might not be able to use.

Therefore an additional document can describe a feedback message, which works as a picture or slice loss indication. If the sender receives this message, it might start with a new intra coded frame, which the receiver would be able to use properly.

The document must specify the format of the FCI-field, one or several RTP payload formats, for which this should be used and a FMT number that is not allocated in all of these formats. A different possibility is to define the FMT number and FCI syntax directly in a video RTP Payload Format.

4.3 Application Layer Feedback Messages

These messages are used to transport application defined data directly from the receiver's to the sender's application. The data that is transported is not identified by the feedback message. Therefore the application must be able to identify the messages payload.

Usually applications define their own set of messages, e.g. NEWPRED messages in MPEG-4 or feedback messages in H.263/Annex N,U. These messages do not need any additional information from the RTCP message and thus this message has the following simple format:

```
 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|E|  FMT  |    PT=AFB     |          length               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  SSRC of packet sender                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  SSRC of media source                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
:                     Application Message                       :
```

:   :

The various fields V, P, SSRC and length are defined in the RTP specification [3]. We summarize the meaning of all of the fields below:

version (V): 2 bits
This field identifies the RTP version. The current version is 2.

early packet (E): 1 bit
This bit is set if the packet is sent early. A detailed description can be found in [8].

padding (P): 1 bit
If set, the padding bit indicates that the packet contains additional padding octets at the end which are not part of the control information but are included in the length field. The last octet of the padding is a count of how many padding octets should be ignored.

feedback message type(FMT): 4 bits
This field identifies the type of the feedback. The packet can be used for different purposes. No type is defined up to now:

0:    forbidden
1-15: reserved

packet type (PT): 8 bits
This is the RTCP packet type which identifies the packet as being an Application Feedback Message.

length: 16 bits
The length of this packet in 32-bit words minus one, including the header and any padding. This conforms with the definition of the length field used in RTCP sender and receiver reports [3].

SSRC of packet sender: 16 bits
The synchronization source identifier for the originator of this packet.

SSRC of media source: 16 bits

The synchronization source identifier of the media source that this feedback is related to.

Application Message (FCI): variable length
This field contains the original application message that should be transported from the receiver to the source. The format is application dependent. The length of this field is variable. If the application data is not four-byte aligned, padding must be added.

4.3.1 Example

An example usage of the Application Feedback Message is NEWPRED for MPEG-4. The NEWPRED messages are defined in the MPEG-4 specification. The MPEG-4 entity at the receiver sends the message, which is transported in the application feedback packet. It is then given to the MPEG-4 entity at the sender, which can identify the information as a NEWPRED message and knows what to do with it.

Normally one NEWPRED message is mapped onto one AFM-RTCP packet, and several messages could be continuously mapped onto one AFM-RTCP packet. In the case several messages are mapped onto one AFM-RTCP packet, padding should only be required on the last individual message. One message SHALL NOT be fragmented into different AFM-RTCP packets.

5. Reaction to Feedback

In the previous sections the feedback messages were defined as well as the timing rules when it is allowed to send these messages. We defined only the feedback messages, because we cannot foresee all purposes for which application designers would want to use the low delay feedback and thus the reaction to the messages is generally left to the application. However this section describes some rules and recommendations of what the sender SHOULD or MAY do in the event of receiving a low delay feedback message.

5.1 Congestion Control

Congestion control for real-time or near real-time traffic, especially for unicast, is not new. Recently a mechanism called TFRC

(TCP-Friendly-Rate-Control), [9], was introduced in the TSV working group for standardization. This mechanism controls the sending bit rate according to monitored network conditions in a TCP friendly way, i.e. in the long-term, the RTP traffic shares the bandwidth fair with TCP connections. Other mechanisms are subject of research everywhere and may be applicable for the use in RTP.

Low delay feedback supports the use of these congestion control algorithms. Due to the frequent feedback messages, it is possible for the sender to monitor the network state closely and therefore it is possible to react to upcoming congestion in time. This minimizes the congestion related packet loss and serves the network stability.

A congestion control algorithm that shares the available bandwidth fair with competing TCP connections, e.g. TFRC, SHOULD be used if the low delay RTP session is transmitted in an best effort environment.

5.2 Reactions to Feedback Messages

As described in section 4 feedback messages are either processed directly by the RTP entity or the application that uses RTP. However the exact behavior is implementation, negotiation and environment dependent. In this section we describe how the sender may react and what he should not do.

At reception of a Transport Layer Feedback Message, the RTP entity evaluates the message and reacts to it. At detection of a packet loss, either by receiving a NACK, missing an ACK or other means, the RTP entity MAY invoke error resilience features such as retransmissions. However if the session is transmitted in an best effort environment, the sender SHOULD NOT increase the sending bit rate at detection of a packet loss. Possible reactions are to retransmit important data, while non-important packets that were scheduled for transmission are discarded in the sender.

At reception of a Payload Type or Application Defined Message the application related part is given to the application, while the attached receiver report (all feedback is sent at least in minimum compound packets) SHOULD be evaluated in the RTP entity and MAY serve as the input for the congestion control algorithm (see previous section for details about congestion control).

6. Security Considerations

Security is not considered in this draft.

7. References

1 Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

2 Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997

3 Schulzrinne, H., Casner, S., Frederick, R. and V. Jacobson, "RTP: A Transport Protocol for real-time applications", RFC 1889, January 1996.

4 Postel, J.,"User Datagram Protocol", STD 6, RFC 768, August 1980.

5 Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.

6 Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.

7 Yano, K., Podolsky, M., McCanne, S., "RTP Profile for RTCP-based

Retransmission Request for Unicast session", IETF Internet Draft, work in progress, July 2000.

8 Wenger, S., Ott, J., "RTCP-based Feedback for Predictive Video Coding", IETF Internet Draft, work in progress, July 2000.

9 Handley, M., Padhye, J., Floyd, S., Widmer, J., "TCP Friendly Rate Control (TFRC): Protocol Specification", IETF Internet Draft. work in progress, November 2000.

10 ISO/IEC 14496-2:1999/Amd.1:2000, "Information technology - Coding of audio-visual objects - Part2: Visual", July 2000.

11 ITU-T Recommendation H.263, "Video Coding for Low Bit Rate Communication," November 2000.

8. Author's Addresses

Shigeru Fukunaga
Oki Electric Industry Co., Ltd.
1-2-27 Shiromi, Chuo-ku, Osaka 540-6025 Japan
Tel. +81 6 6949 5101
Fax. +81 6 6949 5108
Email: fukunaga444@oki.co.jp

Noriyuki Sato
Oki Electric Industry Co., Ltd.
1-2-27 Shiromi, Chuo-ku, Osaka 540-6025 Japan
Tel. +81 6 6949 5101
Fax. +81 6 6949 5108
Email: sato652@oki.co.jp

Koichi Yano
FastForward Networks,
75 Hawthorne St. #601
San Francisco, CA 94105
415.430.2500

Akihiro Miyazaki
Matsushita Electric Industrial Co., Ltd
1006, Kadoma, Kadoma City, Osaka, Japan
Tel. +81-6-6900-9192
Fax. +81-6-6900-9193
Mail akihiro@isl.mei.co.jp

Koichi Hata
Matsushita Electric Industrial Co., Ltd
1006, Kadoma, Kadoma City, Osaka, Japan
Tel. +81-6-6900-9192
Fax. +81-6-6900-9193
Mail hata@isl.mei.co.jp

Rolf Hakenberg
Panasonic European Laboratories GmbH
Monzastr. 4c, 63225 Langen, Germany

Tel. +49-(0)6103-766-162
Fax. +49-(0)6103-766-166
Mail hakenberg@panasonic.de

Carsten Burmeister
Panasonic European Laboratories GmbH
Monzastr. 4c, 63225 Langen, Germany
Tel. +49-(0)6103-766-263
Fax. +49-(0)6103-766-166
Mail burmeister@panasonic.de

Full Copyright Statement