Internet Draft
draft-ietf-avt-evrc-08.txt
October 10, 2001
Expires: April 10, 2002

Adam H. Li
UCLA
Editor

An RTP Payload Format for EVRC Speech

STATUS OF THIS MEMO

ABSTRACT

This document describes the RTP payload format for Enhanced Variable Rate Codec (EVRC) Speech. The packet format supports various formats for different application scenarios. A bundled/interleaved format is included to reduce the effect of packet loss on Speech quality. A non-bundled format is also supported for conversational applications.

Table of Contents

1. Introduction

This document describes how compressed EVRC speech as produced by the EVRC codec [1] may be formatted for use as an RTP payload type. Methods are provided to packetize the codec data frames into RTP packets, in interleaved/bundled and zero-header formats. The sender may choose among various formats the best solutions for different application scenarios based on the network condition, bandwidth restriction, delay requirements, and packet-loss tolerance.

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].

2. Background

The Electronic Industries Association (EIA) & Telecommunications Industry Association (TIA) standard IS-127 [1] defines a speech compression algorithm for use in cdma2000 applications. IS-127, or EVRC, is the emerging speech codec standard for cdma2000.

The EVRC codec [1] compresses each 20 milliseconds of 8000 Hz, 16-bit sampled input speech into one of three different size output frames: Rate 1 (171 bits), Rate 1/2 (80 bits), or Rate 1/8 (16 bits). The codec chooses the output frame rate based on analysis of the input speech and the current operating mode (either normal or one of several reduced rates). For typical speech patterns, this results in

an average output of 4.2 kilobits/second for normal mode and lower
for reduced rate modes.

3. RTP/EVRC Packet Format

The RTP timestamp is in 1/8000 of a second units. The RTP payload data for the EVRC codec MUST be transmitted in packets of one of the following two types.

3.1 Type 1 RTP/EVRC Packet Format

This format is intended for the situation where the sender and the receiver use interleaving/bundling to send one or more codec data frames per packet. The RTP packet for this format 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      RTP Header [3]                           |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
| RR| LLL | NNN |                                               |
+-+-+-+-+-+-+-+-+     one or more ToC entries     +-------------+
|                                                 |             |
+-------------------------------------------------+             |
|                                                               |
|                  one or more codec data frames                |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The RTP header has the expected values as described in [3]. The M bit should be set as specified in the applicable RTP profile, for example, RFC 1890 [4]. Note that RFC 1890 [4] specifies that if the sender does not suppress silence (i.e., sends a frame on every 20 millisecond interval), the M bit will always be zero. When multiple codec data frames are present in a single RTP packet, the timestamp is, as always, that of the oldest data represented in the RTP packet. The assignment of an RTP payload type for this new packet format is outside the scope of this document, and will not be specified here. It is expected that the RTP profile for a particular class of applications will assign a payload type for this encoding, or if that is not done, then a payload type in the dynamic range shall be chosen by the sender.

The first octet of a Type 1 format packet is the Interleave Byte. The bits within the Interleave Byte are specified as follows:

Reserved (RR): 2 bits
   Reserved bits. MUST be set to zero by sender, SHOULD be ignored by receiver.

Interleave Length (LLL): 3 bits
   Indicates the length of interleave. MUST have a value between 0

and 7 inclusive (where a value 0 indicates bundling, a special case of interleaving). See Section 5 and Section 6 for more detailed discussion.

Interleave Index (NNN): 3 bits
Indicates the index within a interleaving group. MUST have a value less than or equal to the value of LLL. Values of NNN greater than the value of LLL are invalid.

The Table of Content field (ToC) contains the index(es) for the codec data frame(s) in the packet. There is one entry for each codec data frame. The detailed formats of the ToC field and codec data frame are specified in Section 4.

More than one codec data frames MAY be included in a single Type 1 RTP/EVRC packet by a sender. Multiple data frames may be included within a Type 1 packet with interleaving/bundling format as described in Section 5 and Section 6.

Since no count is transmitted as part of the RTP payload and the codec data frames have differing lengths, the only way to determine how many codec data frames are present in a Type 1 RTP/EVRC packet is to examine the ToC fields of the packet.

3.2 Type 2 RTP/EVRC Packet Format

The Type 2 RTP/EVRC Packet Format is designed for maximum efficiency and low latency in transmission of the EVRC codec data. Exactly one codec data frame MUST be sent in each Type 2 RTP/EVRC packet. There MUST NOT be ToC field preceding the codec data. The EVRC codec rate for the data frame can be found out at the receiver from the length of the codec data frame, since there is only one codec data frame in each Type 2 packet. The Reduce Rate Signal (See Section 4.1) can not be send in-bound with the Type 2 packets because of the lacking of the ToC field in this type.

Use of the RTP header fields for Type 2 RTP/EVRC Packet Format is the same as described in Section 3.1 for Type 1 RTP/EVRC Packet Format. The detailed formats of the codec data frame are specified in Section 4.

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      RTP Header [3]                           |
+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+=+
|                                                               |
+          ONLY one codec data frame            +-+-+-+-+-+-+-+-+
|                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

3.3 Detection Between the Type 1 and Type 2 Packets

All receivers MUST be able to process both types of packets. The sender MAY choose to use one or both types of packets.

The packets of the two types can be distinguished by using different payload type value for the two packet types at the sender and checking the payload type field in the RTP header at the receiver. The association of payload type number with the packet type is done out-of-band, for example by SDP during the setup of a session.

4. Packet Table of Content Entries and Codec Data Frame Format

4.1 Packet Table of Content entries

For each of the codec data frames in Type 1 packets, there is a corresponding Table of Content (ToC) entry. The ToC entry includes flags that indicates if there are more entries following the current one, if rate reduction on the reverse direction is desired, and the rate of the corresponding codec frame. Type 2 packets MUST NOT have the ToC field, and there is always only one codec data frame in each Type 2 packet.

Each ToC entry is one octet in size. The format of the octet is indicated below:

```
 0 1 2 3 4 5 6 7
+-+-+-+-+-+-+-+-+
|F|D|  frm type |
+-+-+-+-+-+-+-+-+
```

Further Entry Indication (F): 1 bit
   Indicates if there are more ToC entries following the current ToC entry. F = 1 indicates the next octet is another ToC entry. F = 0 indicates that the current entry is the final ToC entry.

Reduce Rate (D): 1 bit
   Setting the 'D' bit indicates that the sender is requesting a reduced codec rate for the reverse direction. When the 'D' bit is not set, the sender is requesting that the codec resume normal operation. In the case of packet loss, the codec SHOULD continue to operate in the mode indicated by the last codec frame received. Receivers are NOT REQUIRED to respond to the Reduce Rate signal. (See more discussion in Section 8.2).

Frame Type: 6 bits
   The frame type values and size of the associated codec data frame are described in the table below:

Value	Rate	Total codec data frame size (in octets)
0	Blank	0
1	1/8	2

3	1/2	10	
4	1	22	
14	Erasure	0	(SHOULD NOT be transmitted by sender)

All values not listed in the above table MUST be considered reserved. Receipt of a ToC entry with a reserved value in Frame Type MUST be considered invalid data.

4.2 The Codec Data Frame

The output of the EVRC codec MUST be converted into codec data frames for inclusion in the RTP payload as specified below.

The codec output data bits as numbered in the standard [1] are packed into octets. The lowest numbered bit (bit 1 for Rate 1, Rate 1/2 and Rate 1/8) is placed in the most significant bit (internet bit 0) of octet 1 of the codec data frame, the second lowest bit is placed in the second most significant bit of the first octet, the third lowest in the third most significant bit of the first octet, and so on. This continues until all of the bits have been placed in the codec data frame. The remaining unused bits of the last octet of the codec data frame MUST be set to zero. Note that this is only applicable to Rate 1 frames (171 bits) as the Rate 1/2 (80 bits) and Rate 1/8 frames (16 bits) fit exactly into a whole number of octets.

Following is a detailed listing showing a Rate 1 EVRC codec output frame converted into a codec data frame:

The codec data frame for a EVRC codec Rate 1 frame is 22-byte long. Bits 1 through 171 from the EVRC codec Rate 1 frame are placed as indicated, with bits marked with "Z" set to zero. EVRC codec Rate 1/8 and Rate 1/2 frames are converted similarly, but do not require zero padding because they align on octet boundaries.

```
                 Rate 1 codec data frame (bytes 0 - 3)

 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|0|
|0|0|0|0|0|0|0|0|0|1|1|1|1|1|1|1|1|1|1|2|2|2|2|2|2|2|2|2|2|3|3|3|
|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|2|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

                 Rate 1 codec data frame (bytes 19 - 21)

 1           1                   1                   1
 4           5                   6                   7
 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 2 3 4 5
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1|1| | | | | |
|4|4|4|4|4|5|5|5|5|5|5|5|5|5|5|6|6|6|6|6|6|6|6|6|6|7|7|Z|Z|Z|Z|Z|
|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| | | | | |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

5. Interleaving Codec Data Frames in Type 1 Packets

As indicated in Section 3.1, more than one codec data frame MAY be included in a single Type 1 packet by a sender. This is accomplished by interleaving/bundling. Interleaving/bundling of codec data frames is signaled by setting the LLL value in the Interleaving Byte to a value between 0 and 7 inclusive.

The special case with the LLL value set to 0 is a reduced and simplified case of interleaving. This is sometimes called bundling, because multiple consecutive codec data frames are included in one RTP packet in this case. The discussions on general interleaving apply to the bundling case with reduced complexity. The bundling case is discussed in detail in Section 6.

Senders MAY support interleaving/bundling. All receivers MUST support interleaving/bundling.

Given a time-ordered sequence of output frames from the EVRC codec numbered 0..n, a bundling value B, and an interleave length L where n = B * (L+1) - 1, the output frames are placed into RTP packets as follows (the values of the fields LLL and NNN are indicated for each RTP packet):

First RTP Packet in Interleave group:
   LLL=L, NNN=0
   Frame 0, Frame L+1, Frame 2(L+1), Frame 3(L+1), ... for a total of B frames

Second RTP Packet in Interleave group:
   LLL=L, NNN=1
   Frame 1, Frame 1+L+1, Frame 1+2(L+1), Frame 1+3(L+1), ... for a total of B frames

This continues to the last RTP packet in the interleave group:

L+1 RTP Packet in Interleave group:
   LLL=L, NNN=L
   Frame L, Frame L+L+1, Frame L+2(L+1), Frame L+3(L+1), ... for a total of B frames

Senders MUST transmit in timestamp-increasing order. Furthermore, within each interleave group, the RTP packets making up the interleave group MUST be transmitted in value-increasing order of the NNN field. While this does not guarantee reduced end-to-end delay on the receiving end, when packets are delivered in order by the underlying transport, delay will be reduced to the minimum possible.

Receivers MAY signal the maximum number of codec data frames (i.e., the maximum acceptable bundling value B) they can handle in a single RTP packet using the OPTIONAL maxptime RTP mode parameter identified in Section 9.2.

Receivers MAY signal the maximum interleave length (i.e., the maximum acceptable LLL value in the Interleaving Byte) they will accept using the OPTIONAL maxinterleave RTP mode parameter identified in Section 9.2.

Additionally, senders have the following restrictions:

- o MUST NOT bundle more codec data frames in a single RTP packet than indicated by maxptime (see Section 9.2) if it is signaled.
- o SHOULD NOT bundle more codec data frames in a single RTP packet than will fit in the MTU of the underlying network. For the purpose of computing the maximum bundling value, all codec data frames MUST be assumed to have the Rate 1 size.
- o Once beginning a session with a given maximum interleaving value set by maxinterleave in Section 9.2, MUST NOT increase the interleaving value (LLL) exceeding the maximum interleaving value that is signaled.
- o MAY change the interleaving value only between interleave groups.

5.1 Finding Interleave Group Boundaries

Given an RTP packet with sequence number S, interleave length (field LLL) L, interleave index value (field NNN) N, and bundling value B, the interleave group consists of RTP packets with sequence numbers from S-N to S-N+L inclusive. (The sequence numbers used here are for illustrative purposes. When wrapping around happens, the sequence numbers need to be adjusted accordingly). In other words, the interleave group always consists of L+1 RTP packets with sequential sequence numbers. The bundling value for all RTP packets in an interleave group MUST be the same.

The receiver determines the expected bundling value for all RTP packets in an interleave group by the number of codec data frames bundled in the first RTP packet of the interleave group received. Note that this may not be the first RTP packet of the interleave group sent if packets are delivered out of order by the underlying transport.

On receipt of an RTP packet in an interleave group with other than the expected bundling value, the receiver MAY discard codec data frames off the end of the RTP packet or add erasure codec data frames to the end of the packet in order to manufacture a substitute packet with the expected bundling value. The receiver MAY instead choose to discard the whole interleave group.

5.2 Reconstructing Interleaved Speech

Given an RTP sequence number ordered set of RTP packets in an interleave group numbered 0..L, where L is the interleave length and B is the bundling value, and codec data frames within each RTP packet that are numbered in order from first to last with the numbers 1..B, the original, time-ordered sequence of output frames from the EVRC codec may be reconstructed as follows:

```
First L+1 frames:
   Frame 0 from packet 0 of interleave group
   Frame 0 from packet 1 of interleave group
   And so on up to...
   Frame 0 from packet L of interleave group


Second L+1 frames:
   Frame 1 from packet 0 of interleave group
   Frame 1 from packet 1 of interleave group
   And so on up to...
   Frame 1 from packet L of interleave group

And so on up to...

Bth L+1 frames:
   Frame B from packet 0 of interleave group
   Frame B from packet 1 of interleave group
   And so on up to...
   Frame B from packet L of interleave group
```

5.3 Receiving Invalid Interleaving Values

On receipt of an RTP packet with an invalid value of the LLL or NNN field, the RTP packet MUST be treated as lost by the receiver for the purpose of generating erasure frames as described in Section 7.

5.4 Additional Receiver Responsibilities

Assume that the receiver has begun playing frames from an interleave group. The time has come to play frame x from packet n of the interleave group. Further assume that packet n of the interleave group has not been received. As described in section 7, an erasure frame will be sent to the receiving EVRC codec.

Now, assume that packet n of the interleave group arrives before frame x+1 of that packet is needed. Receivers SHOULD use frame x+1 of the newly received packet n rather than substituting an erasure frame. In other words, just because packet n was not available the first time it was needed to reconstruct the interleaved speech, the receiver SHOULD NOT assume it is not available when it is subsequently needed for interleaved speech reconstruction.

6. Bundling Codec Data Frames in Type 1 Packets

As discussed in Section 5, the bundling of codec data frames is a special reduced case of interleaving with LLL value in the Interleave Byte set to 0.

Bundling codec data frames indicates multiple data frames are included consecutively in a packet, because the interleaving length (LLL) is 0. The interleaving group is thus reduced to a single RTP packet, and the reconstruction of the code data frames from RTP packets becomes a much simpler process.

Furthermore, the additional restriction on the senders are reduced to:

- o MUST NOT bundle more codec data frames in a single RTP packet than indicated by maxptime (see Section 9.2) if it is signaled.

- o SHOULD NOT bundle more codec data frames in a single RTP packet than will fit in the MTU of the underlying network. For the purpose of computing the maximum bundling value, all codec data frames MUST be assumed to have the Rate 1 size.

7. Handling Lost RTP Packets

The EVRC codec supports the notion of erasure frames. These are frames that for whatever reason are not available. When reconstructing or playing back speech, erasure frames MUST be fed to the receiving EVRC codec for all of the missing packets.

Receivers MUST use the timestamp clock to determine how many codec data frames are missing. Each codec data frame advances the timestamp clock exactly 160 counts.

Since the interleaving length/bundling value may vary, the timestamp clock is the only reliable way to calculate exactly how many codec data frames are missing when a packet is dropped.

Specifically when reconstructing interleaved speech, a missing RTP packet in the interleave group MUST be treated as containing B erasure codec data frames where B is the bundling value for that interleave group.

8. Implementation Issues

8.1 Interleaving Length

The EVRC codec interpolates the missing speech content when given an erasure frame. However, the best quality is perceived by the listener when erasure frames are not consecutive. This makes interleaving desirable as it increases speech quality when packet loss may occur.

On the other hand, interleaving can greatly increase the end-to-end delay. Where an interactive session is desired, either Type 1 with interleaving length 0 or Type 2 RTP payload types are RECOMMENDED.

When end-to-end delay is not a concern, an interleaving length (field LLL) of 4 or 5 is RECOMMENDED.

The parameters maxptime and maxinterleave at the initial setup of the session guarantees that the receiver can allocate a well-known amount of buffer space at the beginning of the session that will be sufficient for all future reception in that session. Less buffer space may be required at some point in the future if the sender decreases the bundling value or interleaving length, but never more buffer space. This prevents the possibility of the receiver needing to allocate more buffer space (with the possible result that none is available).

8.2 Signaling of Reduce Rate

The Reduce Rate signal requests a reduction of the codec rate on the reverse direction. It is NOT REQUIRED that all implementations react to the Reduce Rate signal. If an implementation does react to the Reduce Rate signal, it MUST be able to process/react to the D bit in Type 1 packets. The Reduce Rate signal SHOULD only be used in one-to-one sessions. In multiparty sessions, all the received Reduce Rate signals MUST be ignored.

In addition, the Reduce Rate signal MAY also be sent through non-RTP means, which is out of the scope of this specification.

9. IANA Considerations

One new MIME sub-type as described in this section is to be registered.

The MIME-name for the EVRC codec is allocated from the IETF tree since EVRC is expected to be a widely used codec for Voice-over-IP applications.

The RTP mode has been described in the previous sections.

9.1 Storage Mode

The storage mode is used for storing speech frames, e.g. as a file or e-mail attachment.

The file begins with a magic number to identify that it is an EVRC file. The magic number for EVRC corresponds to the ASCII character string "#!EVRC\n", i.e., "0x23 0x21 0x45 0x56 0x52 0x43 0x0A" in network byte order.

The codec data frames are stored in consecutive order, with a single

TOC entry field (1 octet) prefixing each codec data frame. The F bit and the D bit in the ToC entry field SHOULD be set to 0 and MUST be ignored when processing speech data from storage mode.

Speech frames lost in transmission and non-received frames MUST be stored as erasure frames (frame type 14, see definition in Section 4.1) to maintain synchronization with the original media.

9.2 EVRC MIME Registration

Media Type Name: audio

Media Subtype Name: EVRC

Required Parameters:

ptype: Indicates the Type of the RTP/EVRC packets. The valid values are 1 (Type 1) or 2 (Type 2).

Optional parameters for RTP mode:

ptime: Defined as usual for RTP audio [5].

maxptime: The maximum amount of media which can be encapsulated in each packet, expressed as time in milliseconds. The time SHALL be calculated as the sum of the time the media present in the packet represents. The time SHOULD be a multiple of the duration of a single codec data frame (20 msec). If not signaled, the default maxptime value SHALL be 200 milliseconds.

maxinterleave: Maximum number for interleaving length (field LLL in the Interleaving Byte). The interleaving lengths used in the entire session MUST NOT exceed this maximum value. If not signaled, the maxinterleave length SHALL be 5.

Optional parameters for storage mode: none

Encoding considerations for RTP mode: see Section 5 and Section 6 of RFC xxxx.

Encoding considerations for storage mode: see Section 9.1 of RFC xxxx.

Security considerations: see Section 11 "Security Considerations" of RFC xxxx.

Public specification: RFC xxxx.

Additional information for storage mode:
Magic number: #!EVRC\n
File extensions: evc, EVC
Macintosh file type code: none

Object identifier or OID: none

Intended usage: COMMON. It is expected that many VoIP applications (as well as mobile applications) will use this type.

Person & email address to contact for further information:
adamli@icsl.ucla.edu

Author/Change controller:
adamli@icsl.ucla.edu
IETF Audio/Video transport working group

10. Mapping to SDP Parameters

Please note that this section applies to the RTP mode only.

Parameters are mapped to SDP [5] as usual.
Example usage in SDP:

```
m = audio 49120 RTP/AVP 97
a = rtpmap:97 EVRC
a = fmtp:97 ptype=1; maxinterleave=2
a = maxptime:80
```

11. Security Considerations

RTP packets using the payload format defined in this specification are subject to the security considerations discussed in the RTP specification [3], and any appropriate profile (for example [4]). This implies that confidentiality of the media streams is achieved by encryption. Because the data compression used with this payload format is applied end-to-end, encryption may be performed after compression so there is no conflict between the two operations.

A potential denial-of-service threat exists for data encoding using compression techniques that have non-uniform receiver-end computational load. The attacker can inject pathological datagrams into the stream which are complex to decode and cause the receiver to become overloaded. However, this encoding does not exhibit any significant non-uniformity.

As with any IP-based protocol, in some circumstances, a receiver may be overloaded simply by the receipt of too many packets, either desired or undesired. Network-layer authentication may be used to discard packets from undesired sources, but the processing cost of the authentication itself may be too high. In a multicast environment, pruning of specific sources may be implemented in future versions of IGMP [6] and in multicast routing protocols to allow a receiver to select which sources are allowed to reach it.

Interleaving MAY affect encryption. Depending on the used encryption scheme there MAY be restrictions on for example the time when keys can be changed.

12. Acknowledgements

The editor thanks the following authors for contributions to this document: J. D. Villasenor, D.S. Park, J.H. Park, K. Miller, S. C. Greer, D. Leon, N. Leung, K. J. McKay, M. Lioy, T. Hiller, P. J. McCann, M. D. Turner, A. Rajkumar, D. Gal, M. Westerlund, L.-E. Jonsson, G. Sherwood, and T. Zeng.

13. References

[1] TIA/EIA/IS-127, "Enhanced Variable Rate Codec, Speech Service Option 3 for Wideband Spread Spectrum Digital Systems", January 1997.

[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] Schulzrinne, H., "RTP Profile for Audio and Video Conferences with Minimal Control", RFC 1890, January 1996.

[5] M. Handley and V. Jacobson, "SDP: Session Description Protocol", RFC 2327, April 1998.

[6] Deering, S., "Host Extensions for IP Multicasting", STD 5, RFC 1112, August 1989.

14. Authors' Address

Adam H. Li
Image Communication Lab
Electrical Engineering Department
University of California
Los Angeles, CA 90095
USA
Phone: +1 310 825 5178
Email: adamli@icsl.ucla.edu

John D. Villasenor
Image Communication Lab
Electrical Engineering Department
University of California
Los Angeles, CA 90095
USA

Phone: +1 310 825 0228
Email: villa@icsl.ucla.edu

Dong-Seek Park
Samsung Electronics
Suwon, Kyungki 442-742
Korea
Phone: +82 31 200 3674
Email: dspark@samsung.com

Jeong-Hoon Park
Samsung Electronics
Suwon, Kyungki 442-742
Korea
Phone: +82 31 200 3747
Email: dspark@samsung.com

Keith Miller
Nokia
6000 Connection Drive
Irving, Texas 75039
USA
Phone: +1 972 894 4296
Email: keith.miller@nokia.com

S. Craig Greer
Nokia
6000 Connection Drive
Irving, Texas 75039
USA
Phone: +1 972 894 4867
Email: craig.greer@nokia.com

David Leon
Nokia
6000 Connection Drive
Irving, Texas 75039
USA
Phone: +1 972 374 1860
Email: david.leon@nokia.com

Marcello Lioy
QUALCOMM, Incorporated
5775 Morehouse Drive
San Diego, CA 92121
USA
Phone: +1 858 651 8220
Email: mlioy@qualcomm.com

Nikolai Leung
QUALCOMM, Incorporated

7710 Takoma Ave.
Takoma Park, MD 20912
USA
Phone: +1 703 346 8351
Email: nleung@qualcomm.com

Kyle J. McKay
QUALCOMM, Incorporated
5775 Morehouse Drive
San Diego, CA 92121-1714
USA
Phone: +1 858 587 1121
EMail: kylem@qualcomm.com

Tom Hiller
Lucent Technologies
263 Shuman Drive, Room 2F-218
Naperville, IL 60137
USA
Phone: +1 630 979 7673
Email: tom.hiller@lucent.com

Peter J. McCann
Lucent Technologies
263 Shuman Drive, Room 2Z-305
Naperville, IL 60137
USA
Phone: +1 630 713 9359
Email: mccap@lucent.com

Michael D. Turner
Lucent Technologies
67 Whippany Rd, Room 2A-203
Whippany, NJ 07981
USA
Phone: +1 973 386 3579
Email: mdturner@lucent.com

Ajay Rajkumar
Lucent Technologies
67 Whippany Rd, Room 1A-235
Whippany, NJ 07981
USA
Phone: +1 973 386 5249
Email: ajayrajkumar@lucent.com

Dan Gal
Lucent Technologies
67 Whippany Rd
Whippany, NJ 07981
USA
Phone: +1 973 428 7734
Email: dgal@lucent.com

Magnus Westerlund
Ericsson Radio Systems AB
Torshamnsgatan 23
SE-164 80 Stockholm
Sweden
Phone: +46 8 4048287
Email: magnus.westerlund@ericsson.com

Lars-Erik Jonsson
Ericsson Erisoft AB
Box 920
SE-971 28 Luleà
Sweden
Phone: +46 920 20 21 07
Email: lars-erik.jonsson@ericsson.com

Greg Sherwood
PacketVideo Corporation
4820 Eastgate Mall
San Diego, CA 92121
USA
Email: sherwood@packetvideo.com

Thomas Zeng
PacketVideo Corporation
4820 Eastgate Mall
San Diego, CA 92121
USA
Email: zeng@packetvideo.com