

# Videoconferencing on the Internet

Thierry Turletti and Christian Huitema

**Abstract**—This paper describes the INRIA Videoconferencing System (IVS), a low bandwidth tool for real-time video between workstations on the Internet using UDP datagrams and the IP multicast extension. The video coder-decoder (codec) is a software implementation of the UIT-T recommendation H.261 originally developed for the Integrated Services Digital Network (ISDN). Our focus in this paper is on adapting this codec for the Internet environment. We propose a packetization scheme, an error control scheme and an output rate control scheme that adapts the image coding process based on network conditions. This work shows that it is possible to maintain videoconferences with reasonable quality across packet-switched networks without requiring special support from the network such as resource reservation or admission control.

## I. INTRODUCTION

AS THE BANDWIDTH available on networks and the speed of computers increases, real-time transmission of video between general purpose work stations becomes a more and more realistic application. However, even with a high speed network, video has to be compressed before transmission. For example, sending uncompressed NTSC video requires about 60 Mb/s. Fortunately, there is so much redundancy in most video sequences that even a relatively simple compression scheme can significantly decrease the rate of video flows. Video compression is generally performed by some form of differential coding, i.e., by sending only the differences between two consecutive images. This leads to highly variable transmission rates because the amount of information to code between two images greatly varies, ranging from very low for still scenes to very high for sequences with many scene changes. Packet switched networks such as the Internet are very well suited for transmitting such variable bit rate traffic [8]. However, videoconferencing requires a minimum level of quality and the Internet does not provide such quality of service (QoS) guarantees yet. Nevertheless, we show that it is possible to obtain good quality using control congestion mechanisms to prevent clobbering of the shared resources.

One can find many video compression algorithms in the literature. Some of them have been standardized such as JPEG [2] for still images, or MPEG [18] and H.261 [25], [19] for moving images. MPEG-1 coding is suited for high definition video storage and retrieval [20]. MPEG-2 extends MPEG-1 to high definition television (HDTV) coding applications [21].

Manuscript received January 3, 1994; revised November 28, 1995 and February 13, 1996; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor S. Ahuja.

The authors are with INRIA, Sophia Antipolis, Cedex, France 06902 (e-mail: turletti@lcs.mit.edu and huitema@bellcore.com).

Publisher Item Identifier S 1063-6692(96)04489-5.

The H.261 standard describes a complex video compression algorithm that allows to achieve a very high compression rate.<sup>1</sup> This standard was designed for use over the Integrated Services Digital Network (ISDN), i.e., for a network with fixed rate channels ( $p \times 64$  kb/s,  $p \in [1, 32]$ ). We have implemented a software version of an H.261 codec for use over the Internet. This implementation is the core of the INRIA Videoconferencing System (IVS) [31]. By adopting a standardized algorithm, IVS can easily interoperate with a large number of H.261-based commercial video codecs [14].<sup>2</sup>

However, this standard is not designed for packet switched networks such as the Internet. Since the Internet does not provide the same QoS as ISDN, we propose a set of schemes to adapt the H.261 video compression algorithm to this environment. In this paper, we describe a packetization scheme, an error control scheme and an output rate control scheme that adapts the image coding process according to network conditions.

These three schemes are developed in Sections III, IV, and V, respectively. Section VI evaluates the performances of IVS. Section VII concludes the paper.

## II. RELATIVE WORK

Without the IP multicast technology [10], the set of videoconferencing tools recently developed in the network research community could never be widely adopted. IP multicast technology extends the traditional IP routing model by providing an efficient multi-party packet delivery. The incremental deployment of IP multicast has been realized through the multicast backbone (MBone), a virtual multicast network built on top the current Internet [6], [26].

IVS is not the only videoconferencing application used by the MBone community. At the same time, when we developed IVS, R. Frederick of Xerox Parc was developing the network video (NV) tool. More recently, S. Mc-Canne at UCB/LBL developed the VIC videoconferencing tool.

NV uses a custom coding scheme tailored for the Internet and targeted for efficient software implementation [11]. Its compression algorithm is based on a Haar Transform, a low computational complexity transform compared to the discrete cosine transform used in H.261. In spite of a lower compression rate performance, NV coding is preferred by the MBone community mainly because of its better run-time performances.

VIC has been built upon the lessons learned from both IVS and NV [27]. It is a flexible application that supports

<sup>1</sup>H.261 video compression rate can be easily adjusted, see Section V-C.

<sup>2</sup>See also, the on-line html document <[http://www.cs.ucl.ac.uk/mice/codec\\_manual/doc.html](http://www.cs.ucl.ac.uk/mice/codec_manual/doc.html)>.

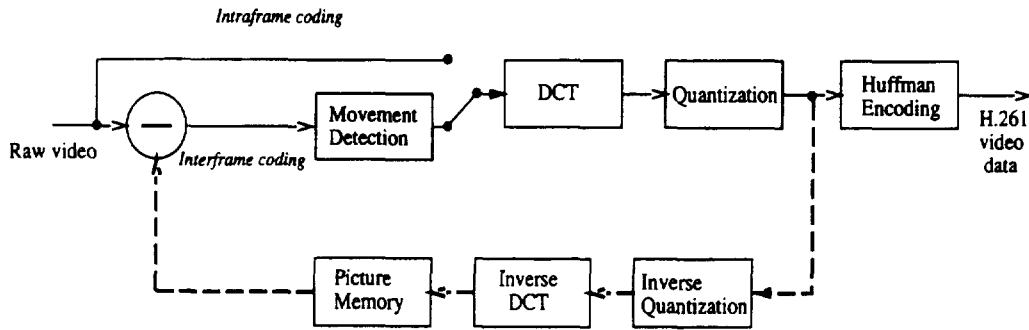


Fig. 1. Basic H.261 coding loop.

multiple network abstractions and several video compression algorithms. VIC can interoperate with both IVS and NV. VIC's H.261 encoder uses only INTRA encoding mode which greatly simplifies the algorithm and improves the run-time performances (in spite of a lower compression rate achieved as shown in Fig. 17).<sup>3</sup>

All these videoconferencing tools are regularly improved and upgraded versions are available in the public domain. Currently, IVS is the only videoconferencing tool that implements a control congestion algorithm by adapting its output rate to the network conditions.

III. H.261 PACKETIZATION

We first give a brief overview of the H.261 video compression standard in order to better understand the following sections.

A. Overview of the ITU-T Recommendation H.261

The H.261 recommendation describes a codec scheme to use for audio-visual services at  $p \times 64$  kb/s ( $p = 1, 2, \dots, 30$ ). An H.261 coder analyzes the successive images of the video stream as sets of blocks of  $8 \times 8$  pixels. The algorithm can be decomposed in several steps: movement detection, discrete cosine transform (DCT), quantization and Huffman encoding.

After performing movement detection, the coder can either decide to encode the difference between the block and its previous encoded/decoded occurrence, or, if there is not enough correlation, to simply encode the new value. This is referred as INTER-frames and INTRA-frames coding, respectively. INTRA-coded codes rely only on the redundancy within a single video frame when inter-coded code also uses the temporal redundancy of video to perform compression. In fact, H.261 does not transmit directly the pixel values or the differences, but rather the coefficients of their DCT [22]. Once transformed, these coefficients are then quantized and Huffman encoded before actual transmission. The coder scheme is shown in Fig. 1.

1) *H.261 Layers*: H.261 coding is organized as a hierarchy of groupings. The video stream is composed of a sequence of images (or pictures) which are themselves organized as a set of groups of blocks (GOB), see Fig. 2. Note that H.261 "pictures" are referred as "frames" in this document. Each GOB holds a set of three lines of 11 macro blocks (MB).

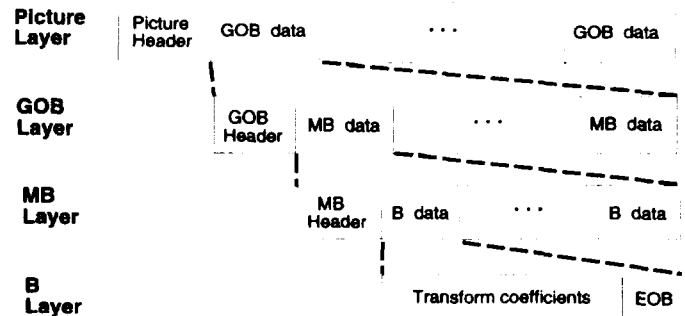


Fig. 2. H.261 layers.

Each MB carries information on a group of  $16 \times 16$  pixels: luminance information is specified for four blocks of  $8 \times 8$  pixels, while chrominance information is given by two red and blue color difference components at a resolution of only  $8 \times 8$  pixels. These components and the codes representing their sampled values are as defined in the ITU-R Recommendation 601 [7].

Two main sizes of images are defined: CIF and Quarter CIF (QCIF).<sup>4</sup> There are 12 GOB's for a CIF picture and three for a QCIF picture as shown in Fig. 3.

This grouping is used to specify information at each level of the hierarchy:

- At the frame level, one specifies information such as the delay versus the previous frame, the image format, and various indicators.
- At the GOB level, one specifies the GOB number and the default quantifier that will be used for the MB's.
- At the MB level, one specifies which blocks are present and which have not changed and, optionally, a quantizer and motion vectors.

B. Considerations for Packetization Over the Internet

H.261 codecs designed for operation over ISDN circuits produce a bit stream composed of several levels of encoding specified by H.261 and companion recommendations. The bits resulting from the Huffman encoding are arranged in 512-bit frames, containing two bits of synchronization, 492 bits of data and 18 bits of error correcting code. The 512-bit frames

<sup>4</sup>CIF is the common interchange format defined by the UIT; it is an interchange format for video images with 288 lines with 352 pixels per line of luminance and 144 lines with 176 pixels per line of chrominance information.

<sup>3</sup>See definition in Section III-A.

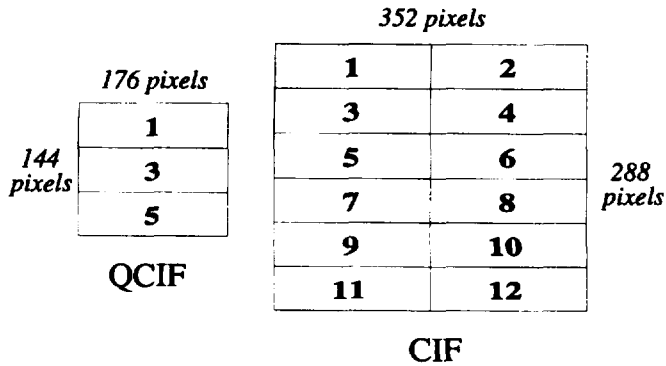


Fig. 3. GOB's arrangement in CIF and QCIF.

are then interlaced with an audio stream and transmitted over  $p \times 64$  kb/s circuits according to specification H.221 [24].

Transmitting a video flow across the Internet requires a different approach. For instance, an application-level error control scheme is more efficient than the 512-bit framing (see Section IV). Similarly, instead of using the H.221 framing, most of multimedia application requirements can be provided by the real time protocol (RTP) (see Section III-C). A comparison between H.261 over ISDN and H.261 over IP is depicted in Fig. 4.

Directly transmitting the result of the Huffman encoding over an unreliable stream of UDP datagrams would have very poor error resistance characteristics. The result of the hierarchical structure of H.261 bit stream is that one needs to receive the information present in the frame header to decode the GOB's, as well as the information present in the GOB header to decode the MB's. However, a video image (or even a GOB) can sometimes be bigger than the maximal transmission unit (MTU).<sup>5</sup> The H.261 recommendation specifies that the maximal size of a CIF image is 32 kbytes (which means 3 kbytes for a GOB, 90 bytes for a MB and 15 bytes for a block). In practice, we observe that the H.261 image size is highly variable according to the quantity of movements and details in the encoded scene: it varies from a few bytes to about twenty kbytes. First versions of the H.261 packetization scheme used a GOB unit of fragmentation. To achieve better performances on lossy environment, the latest version of the packetization scheme takes the MB as unit of fragmentation. In the scheme, packets must start and end on an MB boundary, i.e., an MB cannot be split across multiple packets. Multiple MB's can be carried in a single packet when they fit within the maximal packet size allowed. This practice is recommended to reduce the packet send rate and packet overhead.

In the spirit of the application level framing (ALF) model [9], the H.261 packetization scheme allows each packet received at the decoder to be processed independently. To provide an efficient resynchronization in the presence of packet loss, all the information required to decode an MB independently is sent in a specific RTP-H.261 header conjoined to the H.261 data. This header includes the GOB number in effect at the start of the packet, a reference to the previous MB encoded in this GOB, the quantizer value in effect prior

<sup>5</sup>The MTU size depends on the network, e.g., it is 1536 bytes for an Ethernet network and it can be as low as 576 bytes for the Internet.




Fig. 4. H.261 over ISDN versus H.261 over IP.

to the start of this packet and the reference motion vector data (MVD) for computing the true MVD's contained within this packet.

Moreover, since the compressed MB may not fill an integer number of octets, the H.261 header contains two three-bit integers, SBIT and EBIT, to indicate the number of unused bits in the first and last octets of the H.261 data, respectively.

### C. Overview of RTP

The real time protocol (RTP) aims to satisfy the needs of multi-party multimedia applications: source-identifier, content-identifier, timestamp, demultiplexing, etc. [29]. Moreover, RTP allows interoperability between the existing Mbone tools.

The RTP specification describes a very thin transport protocol which is the most often integrated into the application processing rather than being implemented as a separate layer. This is in accordance with the ALF spirit [9]. In the IP protocol stack, RTP is situated on top of UDP (see Fig. 4). As a matter of fact, the RTP specification describes two protocols: the data transfer protocol (RTP) and the control protocol (RTCP).

Each RTP data packet is composed of an RTP header followed by the RTP payload (i.e., the data). The RTP header contains a sequence number, a media-specific timestamp and a synchronization source packet identifier (SSRC). Receivers demultiplex packets using the SSRC which is globally unique within an RTP session.

The RTCP manages control information providing mechanisms for data distribution monitoring, cross-media synchronization and sender identification. RTCP packets are transmitted periodically to all participants in the session and the period is adjusted according to the size of the session. In this way, the RTCP bandwidth is limited in order to avoid the NACK explosion problem.

### D. Specification of the Packetization Scheme

The H.261 information is carried as payload data within the RTP protocol. The following fields of the RTP header are specified:

- The payload type should specify H.261 payload format (see the companion RTP profile document RFC 1890).
- The RTP timestamp encodes the sampling instant corresponding to the first MB contained in the RTP data packet. If a video image occupies more than one packet, the timestamp will be the same on all of those packets. Packets from different video images must have different timestamps so that frames may be distinguished by the




Fig. 5. The H.261 header format.

timestamp. For H.261 video streams, the RTP timestamp is based on a 90 kHz clock: this clock rate is a multiple of the natural H.261 frame rate (i.e., 30 000/1001 or, approximately, 29.97 Hz). That way, the clock is simply incremented by the multiple for each frame time.

- The marker bit of the RTP header is set to one in the last packet of a video frame and, otherwise, must be zero. Thus, it is not necessary to wait for a following packet (which contains the start code that terminates the current frame) to detect that a new frame should be displayed.

The RTP-H.261 header will follow the RTP header and precedes the H.261 data as shown in Fig. 5.

The fields in the RTP-H.261 header have the following meanings:

- *Start bit position (SBIT)*: Number of bits that should be ignored in the first data octet.
- *End bit position (EBIT)*: Number of bits that should be ignored in the last data octet.
- *INTRA-frame encoded data (I)*: Set to one if this stream contains only INTRA-frame coded blocks. Set to zero if this stream may or may not contain INTRA-frame coded blocks.
- *Motion vector flag (V)*: Set to zero if motion vectors are not used in this stream. Set to one if motion vectors may or may not be used in this stream.
- *GOB number (GOBN)*: Encodes the GOB number in effect at the start of the packet. Set to zero if the packet begins with a GOB header.
- *Macro block address predictor (MBAP)*: Encodes the macro-block address predictor (i.e., the last MBA encoded in the previous packet).
- *Quantizer (QUANT)*: Quantizer value in effect prior to the start of this packet. Set to zero if the packet begins with a GOB header.
- *Horizontal motion vector data (HMVD)*: Reference horizontal motion vector data (MVD). Set to zero if V flag is zero or if the packet begins with a GOB header, or when the MTYPE of the last MB encoded in the previous packet was not motion compensated. HMVD is encoded as a two's complement number.
- *Vertical motion vector data (VMVD)*: Reference vertical motion vector data (MVD). Set to zero if V flag is zero or if the packet begins with a GOB header, or when the MTYPE of the last MB encoded in the previous packet was not motion compensated. VMVD is encoded as a two's complement number.

1) *Recommendations for Hardware Codecs*: Packetizers for hardware codecs can trivially figure out GOB boundaries using the GOB-start pattern included in the H.261 data. The

cheapest packetization implementation consists of splitting the video flow at the GOB level by sending an entire number of GOB's in a packet. But when a GOB is too large, the packetizer has to parse it in order to perform MB fragmentation. Note that this requires relatively little processing since it is not necessary to fully decompress the H.261 stream to fill in the H.261 specific header. However, we recommend to use MB level fragmentation when feasible in order to reduce the output packet rate and, therefore, decrease the overhead.

At the receiver, the data stream can be depacketized and directed to a hardware codec's input. If the hardware decoder operates at a fixed bit rate, synchronization may be maintained by inserting the stuffing pattern between MB's (i.e., between packets) when the packet arrival rate is slower than the bit rate.

The packetization scheme described in this section is currently proposed as standard to the audio-video transport working group (AVT-WG) at the Internet Engineering Task Force (IETF) [32].

#### IV. ERROR CONTROL SCHEME

Errors in a video stream require a different form of correction than errors in a normal data stream. Tests transmitting video stream over a standard TCP connection allowed us to transmit data over the Internet without concern of lost or out-of-sequence packets because of TCP reliability [15]. However, retransmission introduces delays which are not acceptable for real-time applications such as videoconferencing. It is more convenient to use UDP and construct application specific reliability services.

On the Internet, most packet losses are due to congestion rather than transmission errors [3], [28]. Alternatively, packets can be delayed or received out of order. This could happen as a result of the routing and flow control in the network. Due to real-time requirements, delayed video packets are considered as lost packets if delay exceeds a maximum delay value.<sup>6</sup> Using UDP, no mechanism is available at the sender to know if a packet has been successfully received. It is up to the application (i.e., coder and decoder) to handle packet loss and re-sequencing of out of order packets.

Each RTP packet includes a header in which a sequence number field and timestamp are stored. The sequence number is incremented by one for each packet sent, whereas, the timestamp reflects the time when the frame was grabbed. Packet losses can be detected using the RTP sequence number.

The H.261 algorithm uses the temporal redundancy of video to perform compression. The point is that differential coding (or INTER coding) is sensitive to packet loss. Fig. 7 shows a classical "head and shoulders" video sequence usually called *Miss America*. Fig. 7(a) shows the effect of packet loss one image after the loss occurred. In this experiment, the *Miss America*'s head was moving to the right. We note that a lot of blocks around the face are corrupted. Actually, the part of the image affected by the loss will remain blurred as long as all corresponding MB's are not INTRA encoded. There are several ways to mitigate packet loss:

<sup>6</sup>The maximum delay value is empirically set to 100 ms in IVS.




Fig. 6. Data and NACK packets.


- The safest way consists of using only INTRA-frame encoding and MB level conditional replenishment.<sup>7</sup> INTRA coding is much less efficient than INTER coding because a large amount of temporal redundancy exists in image sequences. Removing the INTER coding will result in significant decrease of the compression ratio.<sup>8</sup>

- A more efficient way consists of replenishing, in INTRA mode, only the MB's concerned by the loss. As all GOB belonging to a given video image carry the same timestamp, the receiver can determine a list of GOB's really received for that timestamp and, thus, identifying the missing blocks. Requesting a specific re-initialization of these missing blocks through a negative acknowledgment (NACK) packet is more efficient than requesting a complete refreshment of the image through a full INTRA request (FIR) packet. Fig. 6 shows the NACK emission after a packet loss. In this example, the coder uses QCIF and all GOB's are sent in different packets. When the decoder notices that it didn't receive packet B, it sends a NACK packet to the coder. The NACK information means that GOB 3, image  $n$  is lost. The encoder will put all the MB encoded in the lost packet B into packet E, using INTRA mode. Actually, this is a forced replenishment and not a retransmission procedure because encoding occurs for a new frame and not for the previous lost frame. Fig. 7(a) shows the image after the replenishment procedure.<sup>9</sup> However, the NACK-based method can lead to the feedback explosion problem when receivers are too numerous. If all participants generate NACK's packets each time a packet is lost, network congestion problems will appear very soon. Also, regular hardware H.261 codecs are not designed to accept NACK packets.


<sup>7</sup>This method is currently used by the VIC videoconferencing tool.

<sup>8</sup>We estimate that the INTER mode increases the compression ration by about 30%, see Fig. 17.

<sup>9</sup>In this case, the replenishment happened three images after the blurred image (i.e., 300 ms after for this 10 f/s experiment).



(a)



(b)

Fig. 7. INTRA refreshment after NACK receipt.

- A third way consists of periodically refreshing the image in INTRA encoding mode. For control of accumulation of inverse transform mismatch error, the H.261 recommendation requires INTRA encoding of each MB at least once every 132 times it is transmitted. In order to speed the recovery procedure, the coder can adapt the INTRA refreshment rate according to the packet loss rate information.<sup>10</sup>

IVS implements the second and the third methods, and the method is selected according to the size of the session: we empirically set the threshold to 10 participants. When there are fewer than 10 receivers, NACK's packets are used. Else, the INTRA refreshment rate is adapted to the network congestion state. Let us examine what it means in terms of bandwidth. The two extreme cases for multicasting distribution are the star and the chain network topologies, see Fig. 8. In the following analysis,  $p$  is the average packet loss observed

<sup>10</sup>Receivers can periodically send back to the video coder the packet loss rate they observe [5].




Fig. 8. Star and chain network topologies.

by receivers,  $R$  is the number of receivers in the session and  $N$  is the number of data packets sent during the  $T$  interval of time. In the star network, during the  $T$  interval of time,  $(RN)$  packets are sent by the sender and  $(pRN)$  NACK packets are sent by the receivers. The numbers are  $(N)$  and  $(pRN)$  for a chain network, respectively. Then, the proportion of NACK packets to data packets is within the interval  $[p, pR]$ . The corresponding bandwidth proportion must take into account the length of the data and feedback packets. With an average of 500-bytes per data packet and 12 bytes per NACK packet sent, the proportion of bandwidth used by the feedback channel is within the interval  $[\frac{12}{500}p, \frac{12}{500}pR]$ . In IVS, the maximal  $R$  value is 10. For example, if the session has ten participants and if the average packet loss rate is 20%, the corresponding feedback traffic remains below 5% of the data traffic (between 0.48 and 4.8% depending on the network topology).

## V. CONGESTION CONTROL SCHEME

Videoconference on the Internet could well be a *killer application*; the network administrators nightmares are probably already populated by thousands of hosts all sending megabits of videos over the net and swamping the T3 based backbones. The text book solution for controlling video over the network is called "resource reservation," generally combined with "admission control." To put it shortly, whoever wants to start a video transmission is supposed to ask for permission first, requesting a certain amount of resources. It is assumed that the characteristics of the data flow is known before starting the session and that the network has enough knowledge to decide if such a flow should be admitted or not. There is one advantage to this model: once accepted, one has a guaranteed QoS. But there are also many drawbacks, like the need to couple admission with accounting, the need to enforce complex reservation schemes and, last but not least, the need to introduce a virtual circuit philosophy in an otherwise purely datagram network. Intensive work is currently in progress in the IETF to provide Internet applications the QoS required for their data flows (see the *ReSerVation setup Protocol* (RSVP) [34]). However, since such resource reservations are not yet currently deployed, we investigated another solution, trying to validate, for video, the end-to-end control model that had been so successful for classic data exchange.

End-to-end control needs two components: a network sensor and a throughput controller. Feedback mechanisms for video sources have been proposed for networks with variable capacity channels such as the Internet. There, the goal is to adjust the parameters (and, hence, the output rate) of video coders based on feedback information about changing network conditions, i.e., changing capacity in the network. Gilge *et*

*al.* propose to use feedback control, but they do not describe specific control mechanisms [12]. Wakeman at UCL describes a specific scheme [33]. However, this scheme requires that the source of a connection knows the service rate of the bottleneck on this connection. This service rate can be estimated in networks where the switches use a so-called fair queueing or equivalent discipline, for example, using the packet pair mechanism described in [17]. However, it is not available in networks with FCFS switches such as the Internet. Other work [16] describes a feedback control scheme which requires that switches send their buffer occupancies and service rates back to the source. Next, we describe the network sensor and the throughput controller implemented in IVS.<sup>11</sup>

### A. The Network Sensor

The Internet infrastructure does not provide sources of traffic with explicit feedback information about the state of the network (e.g., queue occupancies at the switches). The only easily available information is implicit information such as measures of losses and/or round-trip delays. In IVS, we use a feedback information based on measured packet losses.

In order to monitor how many video packets arrive at their destinations via multicasting, a source should obtain information from each receiver indicating packet loss on the path from the source to that receiver. One possible approach is to let each receiver send a NACK packet whenever it detects a loss. However, this can lead to the well-known NACK explosion when the network is congested. Another approach consists of periodically sending a QoS measure which is the packet loss rate observed by receivers during a time interval of length  $T$ . We refer to  $T$  as the averaging interval. In IVS, we take  $T$  to be the time required at a receiver to get 100 packets.<sup>12</sup> As suggested in the RTP draft document, we also make sure that each receiver sends feedback information at least once every 2 min.

It is clear that the QoS approach is more efficient than the NACK approach as soon as the packet loss rate is higher than 1%, which is almost always the case on the Internet. To further decrease the impact of feedback traffic on the network, each receiver delays its feedback message by a random amount of time drawn from the range  $[0 \cdots T]$ . This technique is employed to prevent receivers from sending back their feedback at the same time which could create periodic congestion on the network.

Each receiver sends its measured loss rate back to the source using RTP.<sup>13</sup> Then the source converts the different measures of QoS into a global measure characterizing how well packets arrive at the receivers. Our approach is to use the median loss rate.

### B. Throughput Controller

Control actions are taken by the coder at discrete points in time, specifically, whenever a sequence of 100 packets has

<sup>11</sup> A more detailed description can be found in [4].

<sup>12</sup> Observe that interval lengths might be slightly different for different receivers.

<sup>13</sup> The RTP specification provides a framework to send QoS information [29].

been encoded and sent. Of course, the number 100 is chosen so that the interval between successive controls is approximately equal to the interval over which the feedback information is computed at the receivers.

During a control action, the control algorithm adjusts the maximum output rate of the coder  $max\_rate$  so that the median loss rate stays below a tolerable loss rate. The median loss rate is denoted by  $med\_loss$ , and the tolerable loss rate by  $tol\_loss$ . Specifically,  $max\_rate$  is decreased by a factor of two if the median loss rate is larger than  $tol\_loss$ . Otherwise, it is increased by a fixed fraction of its current value. We also make sure that the output rate is always larger than some minimum rate to guarantee a minimum quality of the videoconference at the receivers. Receivers whose connections have an insufficient quality are expected to either obtain more resource through some reservation mechanism, or leave the conference.

Thus, the control algorithm is as follows:

```

if ( $med\_loss > tol\_loss$ )
   $max\_rate = \max(max\_rate/2, min\_rate)$ ;
else
   $max\_rate = gain * max\_rate$ .

```

In IVS, we set  $min\_rate = 10$  kb/s,  $gain = 1.5$ , and  $tol\_loss = 10\%$ . We also set the maximum value of  $max\_rate$  to 100 kb/s. In these experiments, the video source is an IVS source at INRIA. The number of receivers is such that the environment is a large multicast environment, i.e., receivers send QoS packets periodically back to the source. We analyze the connection between the IVS source at INRIA Sophia Antipolis and a receiver at University College London (UCL). Note, however, that the connection between UCL and INRIA is a multicast connection, i.e., the packets sent over the connection are carried over the Mbone. At this time, the multicast path from France to Great Britain goes through CERN in Geneva and Amsterdam in the Netherlands.<sup>14</sup>

Figure 9 shows the evolution of the maximum output rate  $max\_rate$  at the source (plain line) and the average packet loss computed at the receiver (dashed line) during 20 min. The unit on the  $x$ -axis is the frame number, the average frame rate was about four images per second.

As expected, we observe that the value of  $max\_rate$  at the source decreases as losses are detected at the receiver. When the packet loss rate is higher than 10%, the  $max\_rate$  value is decreased by half. Then, the IVS video source is capable of adapting its output rate to the network conditions observed.

We have described above the control algorithm used in IVS, i.e., how the maximum output rate of the coder is adjusted based on the feedback information. Next, we describe how the output rate can be adjusted in H.261 codecs and the way it is implemented in IVS.

### C. Output Rate Control for H.261 Codec

Several parameters of a H.261 coder can be adjusted to change the output rate of the coder. The easiest method consists of modifying the frame rate of the application. However,




Fig. 9. Evolutions of  $max\_rate$  (in kb/s) and the loss rate at the receiver (in %).

this solution is only appropriate when the refreshment rate is not a key parameter of the application: experience shows that the rendition of movement is often a very important requirement.

A second way to control the output of the coder is to adjust the value of the quantizer. By using a looser quantization factor for the coefficients, one reduces the precision of the image—this is approximately equivalent to reducing the number of bits per pixel. The resulting coefficients are less variable than the original values, and result in fewer encoding bits after image compression. However, when the quantizer value is set too high, the image becomes blurred during these changes.

Figure 12 shows that the output rate of the coder as a function of the quantizer. These results have been obtained for the well-known “Miss America” test-sequence. When the value of the quantizer decreases, the output rate of the coder increases and so does the image quality (see Fig. 5).


A third way to reduce the data rate is to simply reduce the number of blocks encoded for each image. This can be done by raising the movement detection threshold. This threshold controls the number of the blocks sufficiently different from the previous frame. If the threshold value increases, then the number of blocks to process decreases and, hence, the encoding time and the number of bytes required to encode each image decreases. Increasing the threshold decreases the sensitivity of the coder to movement and, hence, yields a lower quality image.

Within IVS, the user can select two different modes: *privilege quality* and *privilege frame rate* modes. The mode characterizes what parameters are adjusted in the coder so that the output rate stays below  $max\_rate$ .<sup>15</sup>


Privilege quality (PQ) mode is convenient for applications that require high precision in the rendition of individual images (e.g., to transmit slides or still images). In this mode, the values of the quantizer and the movement detection threshold are constant and match the maximal visual quality. Then the coder waits a sufficient amount of time before encoding

<sup>14</sup>This experiment was made in October 1993.

<sup>15</sup>The  $max\_rate$  value is adjustable on the fly by the control congestion algorithm.



(a)



(b)

Fig. 10. CIF image encoded with (a) quantizer values 3 and (b) 7.

the following image so that the output rate stays below the  $max\_rate$  value.

Privilege frame rate (PFR) mode is convenient when the perception of movement is an important factor of quality. The output rate is controlled using different quantizer and movement detection threshold values. We have decided to couple the quantizer with the movement detection threshold using empirical setups.

Actually, it is legitimate to couple these two variables since when the quantizer increases, the precision of the rendition decreases, and the likeliness of a large difference between two frames increases. The state of the codec, in PFR mode, is characterized by the target data rate  $max\_rate$  and a couple of quantizer and detector values: several pairs (quantizer, threshold) have been pre-selected in order to have a linear variation of the output rate. This selection has been obtained by experimentation restricting the quantizer between three and 13, and the threshold between 10 and 35.




Fig. 11. CIF image encoded with a quantizer value of 11.




Fig. 12. Output rate (kb/s) versus frame for a same video sequence using different quantizers.

Since the output rate is rapidly varying, hysteresis is required to prevent undesirable oscillations in the control loop. If the instantaneous rate measured is included in this band, the coupling is preserved. Experiments show that hysteresis of 30% of the maximum bandwidth damps down the output rate. If the instantaneous rate is outside this band, then a new couple is chosen according to the difference between the instantaneous rate and the maximum bandwidth allowed.

The following diagrams have been obtained using a pre-digitized sequence of 200 frames, with QCIF encoding format and three different values of the bandwidth. The first quarter of the sequence is more animated than the rest of the sequence.

Figures 13–15 show the instantaneous rate and the quantizer with  $max\_rate = 10$  kb/s, 30 kb/s, and 50 kb/s, respectively.




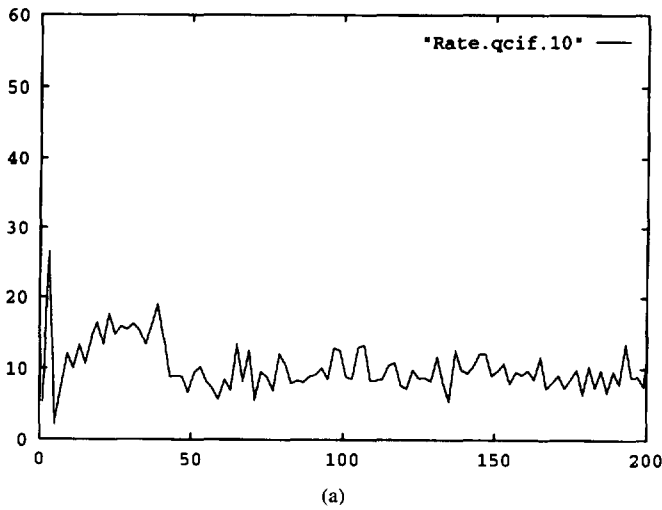


Fig. 13. Output rate and quantizer value versus frame number for  $max\_rate = 10$  kb/s.

Fig. 16 shows the signal-to-noise ratio (SNR) obtained for these three experiments.

The SNR is a mean-square error measure expressed in dB as shown in (1).

$$SNR = -10 \log(MSE), \quad \text{with:}$$

$$MSE = \frac{1}{JK} \sum_{j=1}^J \sum_{k=1}^K [G(j,k) - \hat{G}(j,k)]^2 \quad (1)$$

$$A^2$$

where  $G(j,k)$  denotes the original luminance value of the pixel  $(j,k)$ ,  $\hat{G}(j,k)$  denotes the luminance value of this pixel after encoding/decoding and  $A$  denotes the maximum value of  $G(j,k)$ .

Note that the quantizer selected is inversely proportional to the output rate and that the quality of the image (SNR) is inversely proportional to the quantizer selected. Note also that the quantizer selected is larger during the first quarter of the experiment because the video sequence is more animated, requiring more information to encode.





Fig. 14. Output rate and quantizer value versus frame number for  $max\_rate = 30$  kb/s.



## VI. PERFORMANCE

The output data flow generated by the H.261 coder is nonstationary and rapidly varying. It depends on the quality of the video camera and the type of the images being encoded, as characterized by the number of scene changes, the scene structure, the scene lighting, etc. It also depends on the definition of the images: CIF or QCIF. We also have to take another element into consideration: the computation power of our codecs. The following experiments have been made on a SPARC IPX workstation, with coding and decoding processes running on the same physical machine:

Format	QCIF		CIF	
	Unit	(f/s)	(f/s)	(kb/s)
Expt. 1		6.0	3.6	1.5
Expt. 2		4.0	2.0	16
Expt. 3		1.9	0.5	25

The first experiment is with a still image. When there is no movement, only grabbing and movement detection have to be processed. So, the speed limitation is mainly due to the underlying hardware, in our case the VideoPix board attached to the SparcStation. When there is absolutely no movement, we only encode for each frame the picture and GOB headers.

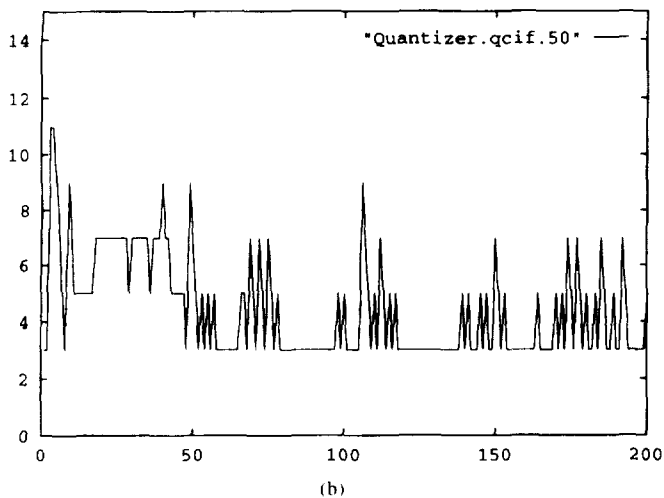
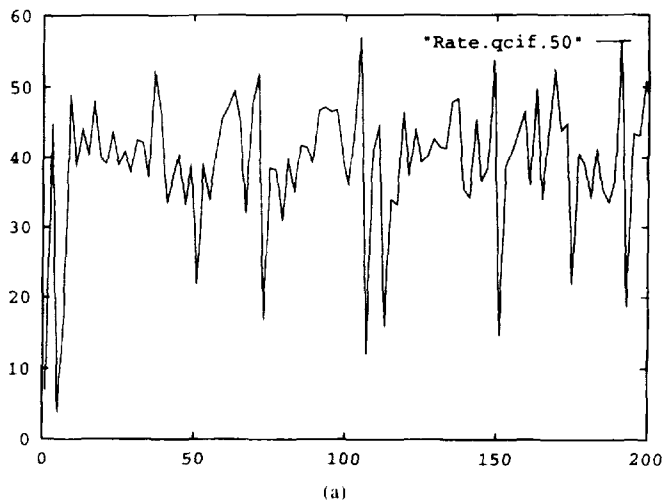


Fig. 15. Output rate and quantizer value versus frame number for  $max\_rate = 50$  kb/s.

The second experiment is characteristic of a classic videoconferencing image, i.e., head and shoulders moving. We can observe that the frame rate is highly dependent on the image size, while the output rate remains almost constant: this is characteristic of a CPU bound application. In fact, the most demanding part of the codec is the computation of the DCT coefficients; the power of the CPU directly limits the number of blocks that can be computed per second, hence, limiting the number of coefficients that have to be transmitted and the number of bits sent on the lines.

The third experiment is for a rapidly changing video scene. In such a case, full INTRA mode encoding is chosen in order to suppress accumulation of errors from INTER mode encoding. INTRA mode encoding usage increases the output rate because more coefficients are encoded by blocks. Image rate decreases because all the blocks have to be encoded and more CPU processing is necessary; the data rate increases because the coefficients are much more dispersed than with differential coding, which makes the Huffman compression less efficient.

The frame rate shown in Table VI is low and might not be suitable for high quality video or remote teaching applications. We found that the frame rate depends on both the video




Fig. 16. SNR (dB) versus frame number for the three experiments 10, 30, and 50 kb/s.




Fig. 17. Frame rate (f/s) versus output rate (kb/s).

grabbing board and the cpu speed. Therefore, we would expect much better performance with higher performance machines. To illustrate this point, we measured the performance of IVS on a SPARCstation 20/501 (bi-processor  $2 \times 75$  MHz) with the VigmaPix board.<sup>16</sup> On this platform, version 3.5 of IVS is able to encode/decode between 25 and 30 fps in QCIF and between 12 and 30 fps in CIF.

Figure 17 shows the performance obtained both for IVS, NV, and VIC on an SS 10/20 (41 Mips) platform with the SunVideo board.<sup>17</sup> The video input sequence is very animated (i.e., all the MB's are encoded in each frame) and QCIF color video encoding is selected without video decoding neither local display functions. We used version 2.6 of VIC to encode both NV and VIC-*H.261* modes and version 3.5 of IVS.<sup>18</sup>

The Sunvideo board allows grabbing up to 20 QCIF frames per second on this platform. Note that neither IVS nor VIC can reach this frame rate when the video sequence is very animated. This is due to the high cpu power required for

<sup>16</sup>See URL <<http://www.vigma.com/products/vigrapix.announce.html>>.

<sup>17</sup>See URL <<http://www.Sun.COM/cgi-bin/show?products-n-solutions/hw/wstns/SunVideo.html>>.

<sup>18</sup>In this experiment, IVS is used with the automatic output rate control disabled.

the H.261 compression algorithm. On the other hand, NV allows obtaining a higher frame rate (in spite of a lower compression rate) because it uses a low computational complexity algorithm. Note that the IVS-H.261 compression rate is about 30% higher than VIC-H.261 compression rate with the same ( $Q = 3$ ) quantizer. This is because IVS uses the INTER encoding mode on top of the INTRA encoding mode. However, the IVS-H.261 coder is more greedy of cpu than the VIC-H.261 [only 8.5 frames can be encoded per second instead of 13 for the same quantizer ( $Q = 3$ )]. Finally, we note that a ( $Q = 13$ ) quantizer gives a compression rate about four times higher than a ( $Q = 3$ ) quantizer.

## VII. CONCLUSION

In this paper, we described a videoconferencing software application for the Internet. This application, available in the public domain<sup>19</sup> is used over the Internet to hold videoconferences, multicast conferences (e.g., the 4th Joint European Networking conference held in Trondheim, Norway, in May 1993), and weekly Multimedia International Conferencing for European Researchers (MICE) seminars [14].<sup>20</sup>

Its main assets are the low bandwidth required, the compatibility with hardware codecs, and the new dimension it brings to classic workstations without high cost, since minimal hardware is necessary. The low bandwidth requirement is very attractive: for instance, low-quality video can be sent on a 9600 b/s link. Moreover, the feedback mechanism we described in this paper allows IVS to behave as a good network citizen.

Further work is ongoing to improve the congestion control algorithm for a heterogeneous multicast environment. Within the Internet, the bandwidth available between several sender-receiver paths can be slightly different. Video gateways can be used to provide different levels of video quality: an application-level gateway for transcoding hardware Motion-JPEG to H.261 video flow has been recently implemented [1].<sup>21</sup> However, a smarter solution to the problem of multicast video transmission over heterogeneous networks consists of using a hierarchical video coding scheme. In such a scheme, the video is split in several flows: the base flow includes the low resolution information, whereas, the enhancement information is sent in additional flows. The idea is to transmit the base flow to all the receivers in the session but to transmit the additional flows only to uncongested receivers. This method will be efficient on the Internet when we are able to associate a higher priority to the essential base flow [30].<sup>22</sup> We are currently working on a wavelet-based video coding scheme which we expect to include in IVS.

<sup>19</sup>IVS sources and binaries are available by [http://zenon.inria.fr/rodeo/ivs/last\\_version](http://zenon.inria.fr/rodeo/ivs/last_version). See also, URL <http://www.inria.fr/rodeo/ivs.html>.

<sup>20</sup>MICE is a European project, aimed at providing appropriate multimedia, multi-party conferencing to researchers in Europe. See also URL <http://www.cs.ucl.ac.uk/mice/mice.html>.

<sup>21</sup>Motion-JPEG stands for JPEG coding applied to moving images. It is based on INTRA coding without movement detection.

<sup>22</sup>Next IP generation (IPng) specifies a priority mechanism associated to the packets sent from a same source using the TCLASS field [23].

## ACKNOWLEDGMENT

The authors gratefully acknowledge S. McCanne, S. Casner and M. Handley for providing helpful comments on the H.261 packetization scheme. J. Bolot contributed to improve the congestion control scheme described in this paper. We thank the MICE research team and the Mbone community who have kindly tested the IVS implementation, reported bugs and provided support for new platforms. Finally, we would like to thank Henry Houh and the anonymous reviewers for their constructive feedback.

## REFERENCES

- [1] E. Amir, S. McCanne, and H. Zhang, "An application-level video gateway," in *Proc. ACM Multimedia '95*, San Francisco, CA, Nov. 1995.
- [2] R. Aravind *et al.*, "Image and video coding standards," *AT&T Tech. J.*, pp. 67-89, Jan/Feb. 1993.
- [3] J. C. Bolot, "End-to-end packet delay and loss behavior in the internet," in *Proc. ACM SIGCOMM'93*, San Francisco, CA, Sept. 1993, pp. 289-298.
- [4] J. C. Bolot, and T. Turletti, "A rate control mechanism for packet video in the internet," in *Proc. IEEE INFOCOM '94*, Toronto, Canada, pp. 1216-1223.
- [5] J. C. Bolot, T. Turletti, and I. Wakeman, "Scalable feedback control for multicast video distribution in the internet," in *Proc. ACM SIGCOMM'94*, vol. 24, no. 4, Oct. 1994, pp. 58-67.
- [6] S. Casner and S. Deering, "First IETF internet audiocast," *ACM CCR*, vol. 22, no. 3, July 1992.
- [7] "CCIR recommendation 601-2: Encoding parameters of digital television for studios," *Int. Telecommun. Union*, 1990.
- [8] D. Clark, "The design philosophy of the darpa internet protocols," in *Proc. SIGCOMM'88 Symp.*, Stanford, CA, Aug. 16-19, 1988, pp. 106-114.
- [9] D. Clark, and D. L. Tennenhouse, "Architectural considerations for a new generation of protocols," in *Proc. ACM SIGCOMM'90*, Sept. 1990, Philadelphia, PA, pp. 200-208.
- [10] S. Deering, "Multicast routing in a datagram internetwork," Ph.D. dissertation, Stanford Univ., CA, Dec. 1991.
- [11] R. Frederick, "Experiences with real-time software video compression," in *Proc. 6th Int. Workshop Packet Video*, Portland, OR, Sept. 1994, pp. F1.1-1.4.
- [12] M. Gilje and R. Gusella, "Motion video coding for packet-switching networks—An integrated approach," in *Proc. SPIE Conf. Visual Commun. Image Processing*, Boston, MA, Nov. 1991.
- [13] J. Guichard, G. Eude, and N. Texier, "State of the art in picture coding for low bitrate applications," *Proc. ICC'90*, pp. 120-125.
- [14] M. Handley, P. T. Kirstein, and M. A. Sasse, "Multimedia integrated conferencing for european researchers (MICE): piloting activities and the conference management and multiplexing centre," *Comput. Networks ISDN Syst.*, pp. 275-290, vol. 26, no. 3, Nov. 1993.
- [15] C. Huitema and T. Turletti, "Software codecs and work station video conferences," in *Proc. INET'92*, Kobe, Japan, pp. 501-508.
- [16] H. Kanakia, P. Mishra, and A. Reibman, "An adaptive congestion control scheme for real-time packet video transport," in *Proc. ACM SIGCOMM'93*, San Francisco, CA, Sept. 1993, pp. 20-31.
- [17] S. Keshav, "A control-theoretic approach to flow control," in *Proc. ACM SIGCOMM'91*, Zurich, Switzerland, Aug. 1991, pp. 1-11.
- [18] D. J. LeGall, "MPEG: A video compression standard for multimedia applications," *Commun. ACM*, no. 4, Apr. 1991.
- [19] M. Liou, "Overview of the  $p \times 64$  kb/s video coding standard," *Commun. ACM*, no. 4, Apr. 1991.
- [20] "Coding of moving pictures and associated audio (MPEG)," ISO/IEC JTC1 SC29.
- [21] "MPEG 2 Video standard," ISO/IEC 13818-2.
- [22] K. R. Rao and P. Yip, *Discrete Cosine Transform: Algorithms, Advantages, Applications*. New York: Academic, 1990.
- [23] S. Bradner and A. Mankin, "The recommendation for the IP Next Generation protocol," *RFC 1752*, Jan. 1995.
- [24] " $p \times 64$  Channel frame structure for audio/video conferencing," ITU-T Recommendation H.221, 1993.
- [25] "Video codec for audiovisual services at  $p \times 64$  kb/s," ITU-T Recommendation H.261, 1993.
- [26] M. R. Macedonia and D. P. Brutzman, "Mbone provides audio and video across the Internet," *IEEE Comput. Mag.*, pp. 30-36, Apr. 1994.

- [27] S. McCanne and V. Jacobson, "Vic: a flexible framework for packet video," *ACM Multimedia*, pp. 511-522, San Francisco, CA, Nov. 1995.
- [28] D. Sanghi, A. K. Agrawala, and B. Jain, "Experimental assessment of end-to-end behavior on internet," in *Proc. IEEE INFOCOM'93*, San Francisco, CA, Mar. 1993, pp. 867-874.
- [29] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RTP: A transport protocol for real-time applications," *RFC 1889*, Nov. 1995.
- [30] T. Turletti, and J-C. Bolot, "Issues with multicast video distribution in heterogeneous packet networks," in *Proc. 6th Int. Workshop Packet Video*, Portland, Oregon, Sept. 26-27, 1994, pp. F3.1-3.4.
- [31] T. Turletti, "The INRIA videoconferencing system (IVS)," *Connexions Mag.*, pp. 20-24, Oct. 1994.
- [32] T. Turletti and C. Huitema, "RTP Payload format for H.261 video streams," *RFC TBD*, 1996.
- [33] I. Wakeman, "Packetized video—Options for interaction between the user, the network and the codec," *The Comput. J.*, vol. 36, no. 1, 1993.
- [34] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, "RSVP: a new resource reservation protocol," in *Proc. IEEE Network*, Sept. 1993.



**Christian Huitema** graduated from the Ecole Polytechnique, Paris, France, in 1975 and received the Ph.D. degree from the Universite' Paris 6 in 1985.

He is now a Research Scientist at Bellcore, after spending nine years at INRIA in Sophia-Antipolis, France, where he led the research project RODEO, investigating various aspects of high-speed networking.

Dr. Huitema is a trustee of the Internet Society and a former chair of the Internet Architecture Board.



**Thierry Turletti** (ACM'96) received the M.S. and the Ph.D. degrees in computer science both from the University of Nice-Sophia Antipolis, France, in 1990 and 1995, respectively. He pursued his Ph.D. studies in the High-Speed Networking Research Group RODEO at INRIA Sophia Antipolis.

His previous research was focused on designing, implementing, and evaluating multimedia applications over the Internet. Since September 1995, he has been a postdoctoral fellow in the Telemedia, Networks and Systems group at LCS, Massachusetts

Institute of Technology. His main research interests include video transmission, flow and congestion control, and signal processing and wireless networking.