

A practical study of video streaming over IEEE 802.11 wireless networks using DirectShow video encoders

Pablo Piñol, Miguel Martínez, Otoniel López, M.P. Malumbres

Dpto. de Física y Arquitectura de Computadores
Universidad Miguel Hernández de Elche
{pablop,mmrach,otoniel,mels}@umh.es

José Oliver

Dpto. de Informática de
Sistemas y Computadores
Univ. Polit. de Valencia
joliver@disca.upv.es

Abstract

Video streaming over wireless networks is an up-to-date subject of investigation due to the present wide-spreading of wireless networks and the increasing number of multimedia applications in mobile environments. This is a challenging task due to the specific characteristics of this kind of networks (node mobility, time-varying wireless channel, signal collision, etc). In this work we have tested several ready-to-use video encoders in a real wireless environment. Our testbed consists of a IEEE 802.11 wireless network and a mobile robot equipped with an onboard webcam and a wireless network adapter which is remotely teleoperated from a control station. The robot sends a video stream to the control station and the control station sends navigation orders to the robot. We will analyze the overall system behavior, standing out specially those aspects related to video encoders and network response.

1. Introduction

IEEE'S 802.11 [11] standard is being increasingly used throughout corporations worldwide due to its good balance of cost, range, bandwidth and flexibility. The bandwidth set by the standard ranges from 1 to 11 Mbps. The new standard 802.11g operates at higher bandwidth (up to 54 Mbps) maintaining the frequency of operation (2.4 GHz) and along with 802.11e [10], that supplies QoS at MAC layer, form a good base to support multimedia traffic.

Although the main application area of wireless networks is related with user access to existing wired network infrastructure, other applications, like telerobotics can benefit from WLANs technology. The development of teleoperated

systems has gained considerable attention due to the new potential applications, such as remote production monitoring [15], remote exploration and manipulation in inhospitable environments [9], tele-surgery [12], and remote training. Important issues concerning communication channels, random propagation delays, bandwidth limitations, fault-tolerance, synchronization, telepresence, and the stability of the robotic systems involving human operators have all been taken into account in different works across the literature [1][4]. Most of them consider Internet as the interconnection network between telecontrolled systems and control stations [7].

The use of wireless channels for teleoperating mobile robots is not new. Most of the existing wireless-controlled robots use a specific (non-standard) radio-modem for communications between control station and robot. But this wireless channel is usually a point-to-point dedicated wireless link which works under good signal quality environments. Therefore, in that case, the wireless link does not represent a drawback in terms of performance of the overall system.

Other recent works describe implementations of teleoperation systems with standard wireless devices. In [6] authors use a PDA with an IEEE 802.11 wireless interface for telecontrolling a robot, and in [18] the same was done through a WAP connection. However, they do not analyze the behavior and performance of wireless link for the correct operation of this kind of applications.

In IEEE 802.11 networks (WLANs) working in infrastructure mode, the access points (base stations) provide the network connectivity between mobile and wired hosts. WLAN nodes share the medium through a CSMA-CA protocol, so network latency may be significant and

dependent of current traffic load. The time-varying conditions of the wireless channel, roaming processes and certain node mobility patterns, may disturb established communications, increasing packet loss ratio and the average delay and jitter, up to intolerant levels for certain kind of multimedia applications.

On the other hand, under this kind of error-prone environments, delivery of encoded video is a challenging task if a minimum QoS is demanded from applications. Video encoders should be prepared to recover damaged bitstreams employing one or several error resilience tools. Video codecs will play an important role in the overall application performance, being critical the ability of concealing network delivery errors to assure the minimum required video quality.

In this paper, we are going to study the feasibility of real-time telerobotic systems in standard WLAN networks, taking into account the different impairments that may appear during a teleoperation session. For that purpose, firstly we will analyze the end-to-end delays of both video and navigation data flows in our own system prototype, so we can determine its limitations. Then, we will carry out some tests under different network scenarios, measuring network parameters as packet delay, jitter and packet loss ratios, in order to determine the stability of the system. Also we will study and analyze the behavior of different video encoders to decide if they fulfill the minimum application requirements.

The paper is organized as follows: in section 2 we describe the teleoperated framework that will be used in this work. Section 3 shows a detailed study of end-to-end delays in both video and navigation data flows. In section 4 we analyze commercial video codecs and how they recover damaged bitstreams in scenarios where packets are lost. In section 5 we devise some real scenarios to measure the impact of different network conditions in our telecontrolled robot system. Finally, in section 6, some conclusions and future work are drawn.

2. Teleoperated mobile robot system

We have developed a basic real-time telerobotic prototype. In Fig. 1, we show our framework, based on an ER1 mobile robot, from Evolution Robotics [5], equipped with a laptop (Pentium IV

2.5 Ghz / 512 Mb), a 802.11b WLAN network adapter and a webcam (labeled as *Server*). The control station (*Client*) consists of another laptop (Centrino 1.7 GHz) connected to a Fast Ethernet network. We also use an 802.11b/g access point which is connected to our robot system through the wireless link and to the control station through the Ethernet network.

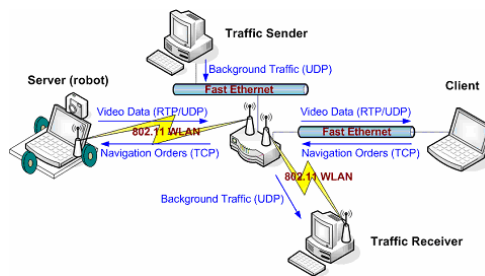


Figure 1. Overview of the telerobotic experimental platform.

The application follows client-server model and is completely based on DirectX/DirectShow architecture. The client software, running at control station, is in charge of requesting a session to the server running at the ER1 robot laptop. This session is mainly based on two data flows: video and navigation.

The first one consists of RTP/UDP (Real-time Transport Protocol) [17] video packets flowing from server to client, so the user at control station can get a video feedback of actual robot location. When the webcam captures a new video frame, the server calls the encoding module for compression. We can choose among several available commercial video codecs or simply choose the uncompressed option. Just behind the encoding module we find the RTP sender block (based on LIVE555 Streaming Media library [14]) where packetization process (fixed size packets) is performed. Then the RTP packets are delivered to the control station (client) through the IEEE 802.11 interface, the access point and the Fast Ethernet network. The control station assembles each frame by arranging all received RTP video packets in one common buffer. When a single encoded video frame is reassembled, the RTP receiver module sends it to the decoding module which recovers the reconstructed version of the original video frame. Finally, the reconstructed

frame is properly rendered on the user interface window at client side.

The navigation flow consists of TCP packets that flow from client to server carrying navigation commands to the robot motion subsystem. When the user issues a navigation command using the GUI (Graphical User Interface), the client sends it to the server (using a short message). When this command arrives at the server, it is immediately processed so the command can reach the robot motion subsystem as fast as possible.

3. End-to-end delay analysis

In this section we are going to analyze the control loop delay sources in both video and navigation flows. For that purpose we have selected three well known test video sequences: *Foreman* video sequence (400 frames) at QCIF format (176x144 pixels), *Coastguard* video sequence (300 frames) at CIF format (352x288 pixels) and *Mobile* video sequence (320 frames) at D1 format (720x756 pixels). For all video sequences we have fixed a common frame rate of 30 fps and a 4:1:1 YUV color space (12 bits/pixel).

We have established three different compression levels: (1) the lowest image quality (a low-but-acceptable quality, around 27 dB of PSNR – *Peak Signal-to-Noise Ratio*), (2) a medium video quality (around 33 dB), and (3) the highest image quality, offering around 39 dB (where further improvements on the quality are nearly imperceptible).

We have chosen three DirectX/DirectShow codecs under evaluation: MPEG-4 [16], H.264 [8] and VP7 [19]. There are two available coding modes: *intra* and *inter*. In *intra* coding mode every video frame is individually encoded. However, when *inter* coding mode is configured, some video frames are individually encoded (*key frames*) and the rest are encoded taking as reference other frames. For *inter* mode we have configured them to produce a key frame every 15 frames (so we label it as *inter-15*).

3.1. Video flow delay chain

Video flow delay is the time elapsed since the robot's webcam captures a single frame until the client renders it to the user. The first source of delay is generated by the robot's webcam. By *vd1*

we represent the time the webcam takes to capture a video frame. From the webcam specifications, it takes from 5 to 20 ms to capture an image. Our webcam has resolutions that range from 160x120 to 640x480 pixels. So, we have lineally determined the capturing delay for each video format.

When compressing video, we will encounter delay *vd2* which measures the time the video encoder takes to compress one video frame. This delay will depend on different variables such as the encoder used, the requested bitrate (opposite to the reconstructed video quality) and other compression parameters. If we do not compress video then *vd2* will have a value of 0. In Table 1 we show the average delay for encoding and decoding a single video frame using a medium level of compression. *Intra* mode is usually faster than *inter* mode but produces a greater number of packets for delivery.

		QCIF - foreman		CIF - coastguard		D1 - mobile	
		Enc.	Dec.	Enc.	Dec.	Enc.	Dec.
H.264	<i>intra</i>	7.48	3.47	27.78	13.13	105.01	42.20
	<i>inter-15</i>	7.60	1.80	35.36	8.20	113.31	17.82
MPEG-4	<i>intra</i>	0.95	1.42	2.73	8.04	9.79	14.90
	<i>inter-15</i>	2.68	0.96	10.43	2.35	36.50	7.81
VP7	<i>intra</i>	14.61	2.43	26.07	8.66	74.53	58.50
	<i>inter-15</i>	24.12	5.50	23.32	14.92	35.98	26.69

Table 1. Encoding and decoding delay (ms).

		QCIF - foreman		CIF - coastguard		D1 - mobile	
		pkt/fr	ms/fr	pkt/fr	ms/fr	pkt/fr	ms/fr
H.264	<i>intra</i>	2.25	0.88	8.94	8.63	28.08	29.70
	<i>inter-15</i>	1.09	0.65	3.75	2.36	5.06	3.91
MPEG-4	<i>intra</i>	2.27	0.80	8.92	7.74	26.91	27.11
	<i>inter-15</i>	1.08	0.51	3.53	2.32	5.08	3.59
VP7	<i>intra</i>	2.22	0.85	8.88	8.25	27.97	29.64
	<i>inter-15</i>	1.07	0.56	3.45	2.30	4.96	3.29

Table 2. Delivery delay (ms) and packets per frame.

By *vd3* we represent the time required to perform the packetization and delivery of one frame. We have used Wireshark [13] network analyzer for measuring network delays. In Table 2 we show this delay using a packet size of 1272 bytes and a medium level of compression. We have measured it without background traffic (no traffic in the IEEE 802.11 and Fast Ethernet segments). We have also measured uncompressed video delivery at a resolution of 160x120 pixels.

Here, the average delivery delay is 24.68 ms (with 23 packets per frame).

Once a packetized frame has been assembled at the client, it is delivered to the uncompressing module to get the reconstructed version of the original frame. The time required by the decoder to reconstruct the original frame is measured by $vd4$. If we do not compress video then $vd4$ is 0. See Table 1 for average decoding delay.

The overall video flow delay will be the sum of all these delays ($vd1 + vd2 + vd3 + vd4$). At a same level of compression (similar bitrates) all codecs produce a similar delay ($vd3$) when delivering a single frame through the network. So the difference between the codecs will be determined by the time needed to encode and decode one frame ($vd2 + vd4$). MPEG-4 has always proved to be faster than the other two codecs.

We have noticed that coding-decoding delay in *inter-15* mode is usually higher than in *intra* mode. At low resolution images (QCIF), the number of packets generated by both methods is very low. So, in the overall video flow delay, using *intra* mode is faster. However, for medium and large resolution images (CIF, D1) the number of packets generated with *intra* mode is significantly greater than *inter-15*. So, in terms of overall video delay, *inter-15* mode is faster than *intra* mode.

3.2. Navigation flow delay chain

The other flow of data carries navigation commands from the client to the server in order to teleoperate the robot. The available navigation commands are the following: move forward/backward, turn left/right, increase/decrease speed, and stop. We send navigation commands in very small TCP packets. We use TCP protocol instead of UDP protocol because it guarantees reliable and in-order delivery of navigation commands.

The first delay we find corresponds with the time elapsed since the user issues a navigation command until the server receives it. We label this delay as $nd1$.

When the server receives a navigation command it immediately sends the order to the robot hardware. By $nd2$ we indicate the time elapsed since the server sends the order until the

robot begins to execute it (and it can proceed to wait for another command). We have measured the average time for executing one navigation command, resulting in the following delays: 22.3, 12.2, and 18.3 ms required by forward/backward, turn and stop commands, respectively.

3.3. Overall delay chain

The overall delay indicates the time the prototype takes for a complete teleoperation action (control loop). It will depend on several factors: frame size, codec used, compression level, navigation command, network status, etc. For example, if we send a forward moving command using an MPEG-4 codec in *intra* mode, a medium level of compression and QCIF video format, we will have an overall delay of 30.78 ms (video delay + navigation command delay). The same configuration but with CIF video format will give an overall delay of 53.31 and 49.9 ms for *intra* and *inter-15* coding modes, respectively.

4. Impact of packet losses on video performance

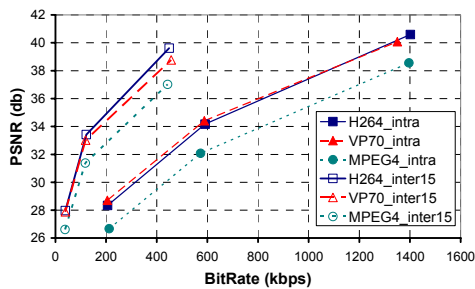
Now, we will analyze the video codec behavior under a scenario with and without packet losses. Bit error patterns are not considered here since IEEE 802.11 network technology uses ACKs for every packet, so the only delivery error is packet loss. In the scenario represented in Fig. 1, the access point (US Robotics USR8054) uses an MTU of 1300 bytes (when working in *turbo mode*) so removing the IP and UDP headers, we will have a maximum RTP packet size of 1272 bytes. We have done our tests with a pair of packet sizes, 1272 and 512 bytes (512 bytes is the default payload size used by some traffic generators). RTP library [14] can be configured to randomly drop packets with a specific PLR (Packet Loss Ratio), so we have used this feature to simulate an error prone environment.

In Fig. 2(b) we show the R/D (Rate/Distortion) performance of the evaluated codecs with *Foreman* (QCIF) video sequence and no packet losses. This figure shows what we obtained in most of our loss-free tests: at similar bitrates H.264/AVC offers the best quality (nearly followed by VP7) and MPEG-4 offers the worst quality with difference. In figure 2(a), we can

perceive the subjective quality difference between the codecs for the first frame of the video sequence, compressed at the lowest bitrate.



(a)



(b)

Figure 2. R/D performance: (a) subjective quality at lowest bitrate, (b) objective PSNR quality comparison.




Figure 3. R/D performance of codecs under packet losses.

In Fig. 3 we evaluate the behavior of the selected video codecs when the video delivery system begins to loose packets. We have used *Foreman* (QCIF) video sequence working in *intra* mode, a medium compression level, and a fixed packet size of 1272 bytes. We show the R/D results with different PLR values. It can be seen that in absence of packet losses H.264/AVC and VP7 get similar results, much better than the ones obtained by MPEG-4. It can also be seen that H.264/AVC does not fall as much as VP7 when the PLR increases.

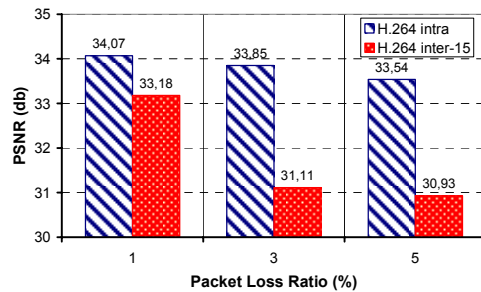


Figure 4. Resilience of *intra* and *inter* compression modes.

Fig. 4 shows the video quality measures for *Foreman* (QCIF) video sequence using H.264/AVC codec with a medium compression level and a packet size of 512 bytes. Here, we show the difference between *intra* and *inter* modes in terms of error resilience. The main reason that explains the lower error resilience behavior of *inter* mode coding is due to the propagation of errors to several frames.

So, if we find an error when decoding one frame, this error will propagate to those frames that require the former one to be reconstructed. As a consequence, the reconstructed video quality will be significantly reduced.

The most outstanding outcomes for the tests we have performed are that, in general, at the same bitrates H.264/AVC and VP7 offer better quality than MPEG-4 (remember that in section 3 we saw that, on the other hand, MPEG-4 was the faster codec) and *intra* mode is more error resilient than *inter* mode.

5. Evaluation under real scenarios

We have performed two test sets under real scenarios. The first one was in an environment with no background traffic, and the second one was in an environment with different levels of background traffic load. All the tests have been done indoor with the robot moving forward at a constant speed of 55 cm/s. It starts moving at the beginning of a corridor nearby the access point (where it has 100% signal strength) and goes away during 110 seconds up to the end of the corridor, where the signal gets lost. Traffic Sender and Traffic Receiver are both static and near the access point (see Fig. 5). We have used *Coastguard* (CIF) video sequence concatenated ten times (so it has a length of 3000 frames and is 100 seconds long) with a frame rate of 30 fps. We have setup the codecs to provide a bitrate of 2 Mbps in *intra* mode.




Figure 5. Signal strength and PLR for different packet sizes.

First, we have done some tests with MPEG-4 codec to evaluate the system performance for different packet sizes without background traffic (unloaded network). In Fig. 5, the packet size that gets the lowest PLR is 1472 bytes (maximum size for an Ethernet MTU of 1500 bytes). So, from now on, we will fix the packet size to 1472 bytes.

The second set of tests was done introducing different background traffic loads. We have used D-ITG [2] traffic generator and D-ITG GUI [3] graphical interface in both sender and receiver to inject background traffic (CBR UDP stream) to our WLAN. For these tests we have used all codecs in *intra* mode with a packet size of 1472 bytes (as mentioned above). We have injected seven different background traffic loads: from

2048 to 3277 Kbps. Traffic loads under 2048 Kbps have a negligible incidence on the video flow delivery and those over 3277 Kbps drive the network to a saturation state.




Figure 6. Jitter comparison at different background traffic loads (MPEG-4 video codec).

In Fig. 6, a comparison of experienced jitter (average packet delay variance) during the session is shown under different background traffic loads. The observed jitter shows a different behavior with respect background traffic (network contention) and received signal strength.

When the signal strength is above 50% (see Fig. 5), the jitter fluctuations are mainly due to network contention (background traffic). However, when signal strength is below 50%, the jitter wildly changes, amplifying the effect that traffic background has over it.

In Fig. 7, we show the time-average PLR in a session where MPEG-4 was used with different background traffic loads (as reference we show the results without background traffic). As it can be seen, packet losses are directly influenced by signal strength and traffic load for all the region of coverage, in a similar fashion than in Fig. 5.

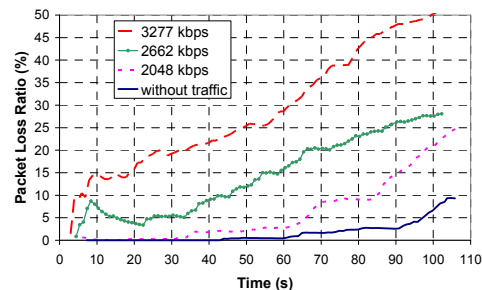


Figure 7. PLR evolution using MPEG-4 codec with different background traffic loads.

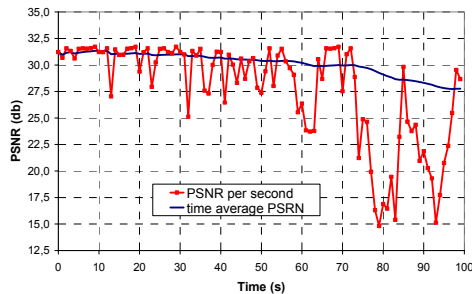


Figure 8. Quality of MPEG-4 reconstructed video with a background traffic load of 2259 Kbps.

In Fig. 8 we show the reconstructed video quality of the received video sequence using the MPEG-4 codec and a background traffic load of 2259 Kbps. We show the evolution of PSNR (each point represents the average PSNR of 30 frames) and the time-average PSNR. It can be seen that when a packet or a burst of packets is lost the PSNR falls drastically (we have found a PSNR decay of up to 15 dB for a single frame). This situation appears in those areas where the received signal strength is under 50%. In those situations there are a lot of lost packets, typically in bursty fashion. So, in most cases, the frame can not be reconstructed (the frame is declared as lost) or, if the codec has enough information for decoding the frame, the quality of the frame recovered is too low.

In Fig. 9 we compare the selected codecs and the video quality they provide with two different traffic loads. As it can be seen, the H.264/AVC video codec always obtains better results than the

other two, showing the same behavior than the one observed without packet losses. However, VP7 shows lower performance than MPEG-4, since this codec is not able to reconstruct a frame when the first part of it is missing. So, to avoid VP7 crashing, we drop the incomplete frame at RTP receiver process. This is the main reason of VP7 low performance.




Figure 9. Quality of reconstructed video with different codecs and traffic loads.

In Fig. 10 we can see the subjective quality of a reconstructed video. Fig. 10(a) shows frame number 2191 of the video sequence without any compression. This corresponds to second 73 of the sequence. In Fig. 10(b) we can see the same frame after it has been encoded and decoded using VP7 codec (without any loss). For this frame we get a value of 31.47 dB of PSNR. In Fig. 10(c) we can see the same frame after reconstruction in a test where packet loss has occurred. This frame offers a value of 14.50 dB of PSNR.



Figure 10. Subjective quality: (a) *Coastguard* original frame number 2191, (b) VP7 encoded-decoded, (c) VP7 reconstruction with packet losses.

6. Conclusions

In this paper we have analyzed the impact of commercial video codecs for video streaming through 802.11 networks running in remotely teleoperated mobile robots testbed. We have developed a client-server application based in DirectX/DirectShow architecture for testing video compression, data delivery and error resilience behavior in a real scenario.

A first analysis was performed to study the teleoperation control loop delays and the performance of video delivery process in order to get a first impression of overall system behavior. For this kind of applications, compression is mandatory for proper operation in order to cope with the minimum video quality and bounded delay demanded by them. The use of video compressors allows us to adjust the required quality without exhausting network resources like available bandwidth and they significantly reduce the control loop delay improving the functionality of this kind of applications.

With respect to packet losses we have observed that H.264/AVC codec is the one that best performance results achieve. Also, we have checked that *intra* coding mode gets better results than *inter* mode, since it avoids error propagation. In the packetization process, the RTP packet size determines the resulting application bandwidth (goodput). Results show that the longer the packet size is, the higher goodput the application gets.

Finally, we think that telerobotic applications running under standard wireless network technologies like IEEE 802.11 can benefit from the use of state-of-the-art video encoders, like H.264/AVC, to improve throughput, error resilience and real-time feedback.

As future work, we plan to use the IEEE 802.11e MAC QoS standard to apply differentiated traffic classes to navigation and video packets in order to improve the application responsiveness.

Acknowledgement

This work was supported by the Spanish Ministry of Science and Technology under grant TIC2003-00339.

References

- [1] Brady, K., Tarn, T.J. Internet-Based Remote Teleoperation. Proc. of the 1998 IEEE Conf. on Robotics and Automation, Belgium, 1998.
- [2] D-ITG, Universita' degli Studi di Napoli, <http://www.grid.unina.it/software/ITG>
- [3] D-ITG GUI, Volker Semken, <http://www.semken.com/projekte>
- [4] Elhajj, I. et al. Supermedia in Internet-Based Telerobotic Operations. IFIP/IEEE Int. Conference on Management of Multimedia Networks and Services, 2001.
- [5] ER1, <http://www.evolution.com/er1>
- [6] Fong, T. et al. Novel interfaces for remote driving: gesture, haptic and PDA. Autonomous Robots, 2001.
- [7] Goldberg, K., et al. Beyond webcams: An introduction to online robots. MIT-press, 2002.
- [8] H.264, <http://x264.nl>
- [9] Hirzinger, G., Brunner, B., Dietrich, J., Heindl, J. Sensor-Based Space Robotics-ROTEX and Its Telerobotic Features. IEEE Trans. on Robotics and Automation, 1993.
- [10] IEEE P802.11 – TG E, MAC Enhancements for Quality of Service, 2002.
- [11] IS.802.11. Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications. IEEE, Inc., August 1999.
- [12] Kwon, D., Woo, K.Y., Cho, H.S. Haptic Control of the Master Hand Controller for a Microsurgical Telerobot System. IEEE Int. Conf. on Robotics and Automation, 1999.
- [13] Lamping, U., Sharpe, R., Warnicke, E. Wireshark user's guide. 2006.
- [14] LIVE555, <http://www.live555.com/livemedia>
- [15] Luo, R., Lee, W.Z., Chou, J.H., Leong, H.T. Tele-Control of Rapid Prototyping Machine Via Internet for Automated Tele-Manufacturing. IEEE Int. Conf. on Robotics and Automation, Vol. 3, pp. 2203-2208, 1999.
- [16] MPEG4, <http://www.free-codecs.com>
- [17] RFC 3550. RTP: A Transport Protocol for Real-Time Applications. 2003.
- [18] Sgouros, N.M., Gerogiannakis, S. Robot teleoperation environments featuring WAP-based wireless devices. Journal of Network and Computer Applications, 2003.
- [19] VP7, <http://www.on2.com/technology/vp7>