

Experimental measurement of vehicular QoS

C. Pinart, I. Barona, C. Alba

Networked Vehicles Division, Telefónica I+D - c/ Emilio Vargas 6, 28043 - Madrid, Spain
{cpg, ibarona, calba}@tid.es

I. INTRODUCTION

Quality of Service (QoS) is broadly defined as “*the collective effect of service performances which determine the degree of satisfaction of a user of the service*” [1]. The prevailing perception in broadband wireless is that as adoption grows, so does the need for guaranteeing QoS. QoS, therefore, has become a critical area of concern for suppliers of broadband wireless access (BWA) equipment and their customers. This is so because QoS directly affects the ability of providers to meet the service level agreement (SLA) requirements demanded by customers. To date, this has been difficult to achieve with most wireless options. Nowadays, vehicular networks and services are becoming a reality, driven by safety requirements and by the investments of car manufacturers and Public Transport Authorities. This means that, in the future, services will be provided in vehicles. These services will be a mixture of existing services such as infotainment or communication, and transport-oriented services, such as safety, driver assistance or driver information services, which will be delivered by broadband vehicular, wireless and cellular networks. As in any network, enforceable QoS will be an essential foundation for widespread acceptance of the mentioned services. In turn, vehicular QoS will allow for more efficient sharing of the operator's infrastructure and, properly managed, will increase subscribe take-up.

For the reasons described above, measuring the QoS of vehicular services is crucial to design future heterogeneous access for the vehicular context and to design suitable SLAs that ensure both customer satisfaction and resource optimization. To this end, Telefónica I+D – the R&D company of the Spanish operator Telefónica - is developing the *Demonstrator for Intelligent Vehicular Environments* (DRIVE) [2], an experimental testbed featuring both advanced vehicular services and vehicular networking, i.e., intra-vehicle, vehicle-to-vehicle (V2V), vehicle-to-roadside (V2R) and vehicle-to-infrastructure (V2I) communications. Thanks to a modular, reconfigurable architecture, DRIVE can integrate different radio interfaces, communications schemes and services easily and readily.

II. THE DRIVE TESTBED

DRIVE consists of a dedicated hardware platform embedded with a software platform containing the communications node, the vehicular services and the human-machine interface (HMI) to be integrated on-board a vehicle. Moreover, DRIVE includes communication and management elements in the network infrastructure. DRIVE's on-board configuration can be used standalone or in a distributed way by setting up connections with other DRIVE-equipped vehicles. Apart from providing proof of concept, DRIVE collects network- and service-related parameters to analyze the performance of vehicular networking and advanced services. Fig. 1 illustrates the generic architecture of the DRIVE demonstrator, where it can be observed that the *Car Gateway* is the core element of the testbed. This gateway integrates the hardware and software components that

enable intra-vehicular, V2V, V2R and V2I communications, advanced services and multimodal HMI. The in-vehicle wireless network considered is composed of wireless sensors located in the vehicle and nomadic devices of the driver and/or passengers. The car gateway may also have access to some of the vehicle sensors and actuators, and has access to the steering wheel and audio system of the vehicle for HMI purposes.

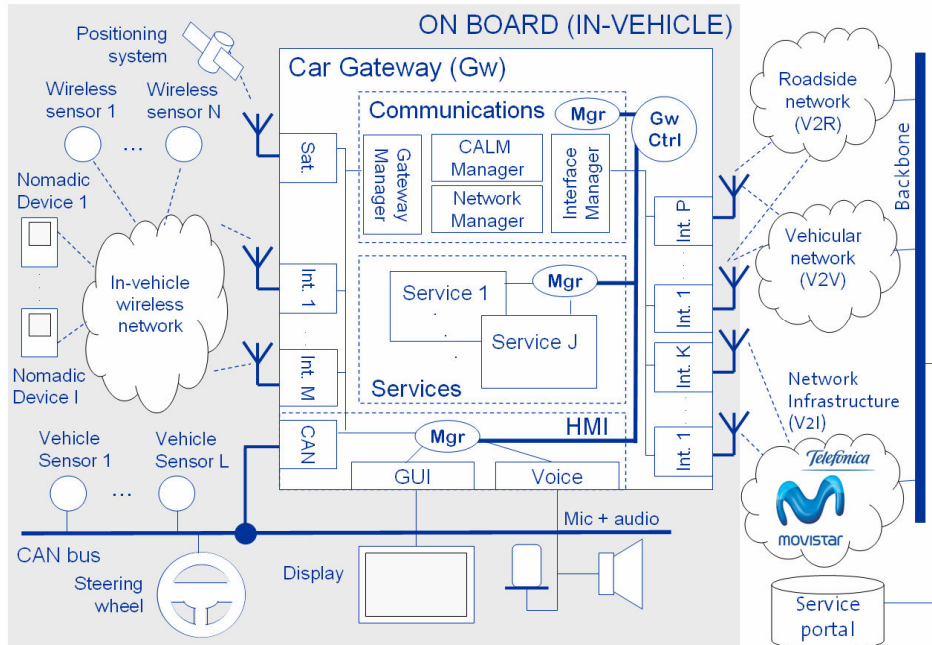


Figure 1. Generic architecture of the DRIVE testbed

As for communications, DRIVE's architecture considers radio interfaces for in-vehicle communications (Int 1 to M in Fig. 1), positioning (Sat in Fig. 1), inter-vehicle and roadside communications (Int 1 to P in Fig. 1) and communications with network infrastructure, such as cellular (Int 1 to K in Fig. 1). To handle all these communications technologies, the car gateway includes an Always Best Connected (ABC) enabled communications node. Fig. 2 illustrates the current implementation of the architecture depicted in Fig.1, which is based on a car-PC architecture with a dual WiFi/3G communications node, in line with common architectures such as that of the European R&D project CVIS (www.cvisproject.org).

III. MEASURING QoS IN DRIVE

There are essentially two ways to provide QoS for network applications: by resource reservation from the network's viewpoint, and by behavior adaptation from the application's viewpoint. In the former, several schemes have been devised which provide QoS in traditional networks. While some schemes reserve network resources during a setup stage [3], other schemes set aside resources on a per-class basis and provide a statistical guarantee of service [4]. Concerning the latter, for an application to achieve resource reservation or behavior adaptation, it first needs to be aware of the network environment. In DRIVE, this is achieved by the communication node embedded in the car gateway (Fig. 1).




Figure 2. DRIVE's current on-board implementation

A. QoS FRAMEWORK FOR SLA VALIDATION

Network-QoS metrics measure connectivity-related aspects such as loss, throughput, delay and delay variation (jitter), sequencing, and total errors. Short-lived sessions typically do not require handover or QoS, while longer-duration sessions (e.g., video streaming), may require handover and typically have QoS requirements including delay, delay variation and packet loss. If we take the example of video streaming in IP-based wireless networks, we can say that no true QoS exists due to challenges such as latency, available bandwidth, variability of wireless link quality, and the high bandwidth demands and variability of video applications. Wireless networks such as IEEE 802.11 are subject to sudden and severe drops in available bandwidth that result in stuttering and frozen video. A common solution to this problem is buffering - although the more variable the wireless link, the bigger the buffer -, which can be combined with rate adaptation. With DRIVE, we are able to measure the relevant quality parameters of critical applications such as video streaming (e.g., latency, errors) to verify if an SLA for these applications could be guaranteed by a network operator in the vehicular context taking into account several constraints from the network (e.g., latency) and the application layers (e.g., buffer size). Examples of such applications are the download of contents (data transfer) and the synchronization with remote multimedia contents for on-the-go playback, e.g., audio.

B. QoS FRAMEWORK FOR ALWAYS-ON CONNECTIVITY

Apart from the challenges inherent to QoS provisioning [3-4], guaranteeing QoS in wireless networks poses additional challenges, namely severe loss, propagation delay, contention, especially in ad-hoc networks, and mobility (e.g., routing, handover). Communications in the DRIVE testbed follow the ABC concept, i.e., always choose the best communication choice possible or available for the given service requirements and user preferences. With ABC, vehicular services will be provided in a multi-technology context involving mobile ad-hoc, wireless and cellular networks, that is, in a context in which diverse but complementary architectures and wireless access technologies will coexist. In general, the proposed interworking architectures (e.g., 3G cellular networks and WiFi) are still hindered by the absence of QoS guarantee [5]. Another important issue in this context is the interworking of QoS at the different layers of the OSI stack. To this respect, [6] shows how

QoS at the IP layer can be interworked with QoS at the wireless link layers to provide the QoS required by the user application. Since QoS is handled differently by each technology, we are defining correspondences between the technologies used in the DRIVE testbed, namely 3G, IEEE 802.11b/g (WiFi family), Bluetooth and Zigbee, and we are deriving service requirements for the management of ABC communications (see example in Fig. 3).




Figure 3. Example of DRIVE's ABC scenario (WiFi and 3G)

IV. SLA VALIDATION OF ON-THE-GO SERVICES

Connectivity (data transport service) between two geographically dispersed locations is usually provided by an independent third party, generically called service provider. The SLA is then a contract agreed upon between the service provider and the service consumer (customer). Originally a financial contract, the SLA evolved towards an instrument for the management of customer's expectations in the late 1990s. Nowadays, an SLA can be defined as a *“contract between a network service provider and a customer that specifies in measurable terms the services furnished by the provider and the penalties assessed if the provider cannot meet the goals established”* [7]. Hence, an SLA details the attributes of the service provided by the service provider to the customer. Examples of attributes are connection uptime, scheduled downtime, and service provider liabilities. End-to-end SLAs are scarce, and among Internet Service Providers (ISP) SLAs are limited, which gives an idea of the non-widespread use of QoS in the 'real world'. The research community has been working on QoS-aware networking infrastructures, such as IETF's Differentiated Services model [4]. However, while there are many papers on QoS mechanisms, only a few describe anything anyone has deployed [8]. And even with a QoS-aware infrastructure deployed and working, there can be an inherent mistrust between a provider and a consumer necessitating a form of SLA with measurable parameters.

Today's wireless ISPs typically include four components in their SLAs:

- **Service availability:** guaranteed percentage of network (connectivity) availability, measured according to a given method (annually, monthly, etc.).
- **Network latency:** maximum round-trip delay (RTD).
- **Packet loss rate:** maximum percentage of lost packets in a given network segment.
- **Throughput:** minimum percentage of average purchased speed (by certified tests).

Table 1 illustrates two examples of real wireless ISP SLAs. Note that SLAs refer to the network part of services, not to the end device used to display and interact with them. It is worth noting that sometimes the perception of users regarding communication services is highly dependant on the terminal used (e.g., cell phone, PDA, handheld device), which has

nothing to do with the wireless ISP, and hence it is not included in the SLA. To the best of our knowledge, today's 3G services have no SLA comparable to those in Table 1.

Provider	Availability	Latency (RTD)	Packet loss	Throughput
Y-net (wireless Internet)	99.99% (annualized)	< 75 ms (backbone) and < 50 ms (last mile)	≤ 1% (backbone)	--
Freewire Broadband (WiMAX)	99.99% (monthly)	≤ 80 ms (Freewire network)	≤ 1% (Freewire network)	≥ 98% average purchased speed

Table 1. SLA examples of wireless Internet Service Providers

A. TEST SETUP

As described in Section III.A, two critical applications for vehicular users are video streaming and download of contents. From the user perspective, web page display is usually a qualitative measure of QoS. In our tests we chose the following services:

1. YouTube-quality video streaming from a Telefónica server. The video has the following characteristics: FLV file with a resolution of 320x240, 25 frames/s and size of 3.2 MB, which are in line with You Tube's average metrics reported in [9] except for the file size, which has been reduced to lower the duration of tests.
2. Large file download (3.2 MB) from the Spanish National Research Network (Red Iris, www.rediris.es).
3. Web page display with large text, CSS and images (data from www.elpais.es) downloaded from a Telefónica server. 70 downloaded files (total of 480 kB).

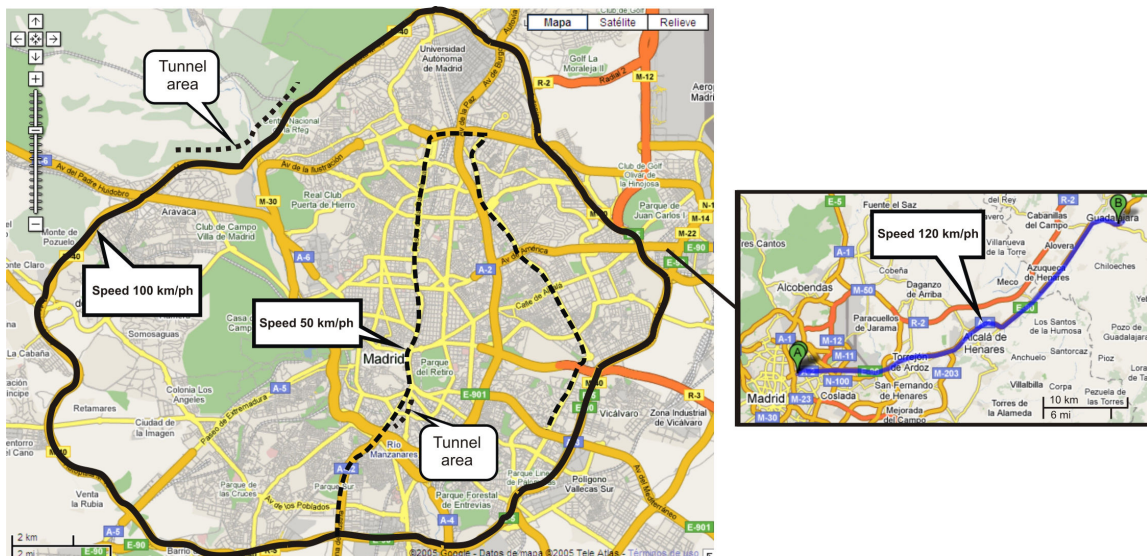


Figure 4. Tested areas in the city of Madrid

For each test, we have two connection strategies: single connection, and per-test (re)connection. Upon a successful connection, each test consists of a suite of four tasks: a 15-packet (64 B each) *ping* is sent to www.rediris.es, and then the services 1, 2 and 3 described above are run. Each test has a limited running time. A test is considered failed if this running time is exceeded. The tests were carried in a semi-urban scenario in the city of Madrid (Fig. 4) at speeds of 50, 100 and 120 km/h, where 200 scenarios were executed.

Due to the limited availability of WiFi in the test areas, 3G was used as the connectivity technology to access the described services from the vehicle. The 3G technology used was High-Speed Downlink Packet Access (HSDPA), which offers a proven performance between 0.92 and 1.48 Mbps [10]. The SLA metrics measured were:

- **Availability**, measured in the number of successful attempts over the total attempts.
- **Latency**, measured as the RTD using the *ping* command to www.rediris.es.
- **Throughput**, measured in regular 1-sec intervals as the number of bits received over each time interval at the network interface of the car gateway assigned to 3G traffic.
- **Packet loss rate (PLR)** was only measured for the ping packets. Additionally, we measured the overhead in the network calculated as the ratio between the total transferred bytes and the expected size of data (e.g., 3.2 MB for the downloadable file). This gives a rough idea of retransmissions, lost packets and protocol overhead.

Service	Availability (%)	Latency (msec)	Overhead (%)	Throughput (bit/sec)
Video streaming	a) 57,30 % b) 46,32 %	a) 334 (average RTD), 1.34%	a) 10,54 b) 10,75	a) 38823,01 b) 35323,80
File download	a) 77,17 % b) 51,02 %	PLR b) 381 (average RTD), 1.48%	a) 3,50 b) 5,80	a) 40523,47 b) 44476,69
Web page display	a) 75,56 % b) 51,55 %	PLR	a) 10,08 b) 9,89	a) 19142,77 b) 18241,00

Table 2. Average QoS measurements a) with and b) without per-test reconnection

B. EXPERIMENTAL RESULTS

Table 2 summarizes the average QoS measurements obtained with the test setup described in the previous section: availability, latency, throughput and network overhead. At the following we comment on the major QoS issues observed during the tests:

- **No impact of speed on the measured throughput:** Fig. 5a illustrates the maximum and average bandwidth as a function of the speed, in 10-km/h intervals. From this figure it can be observed that the speed plays no relevant role in 3G communications. It must be pointed out, though, that the throughput measured at some speed intervals (60-80 and over 110) is lower because of the tunnel areas in the test route (Fig. 5b).

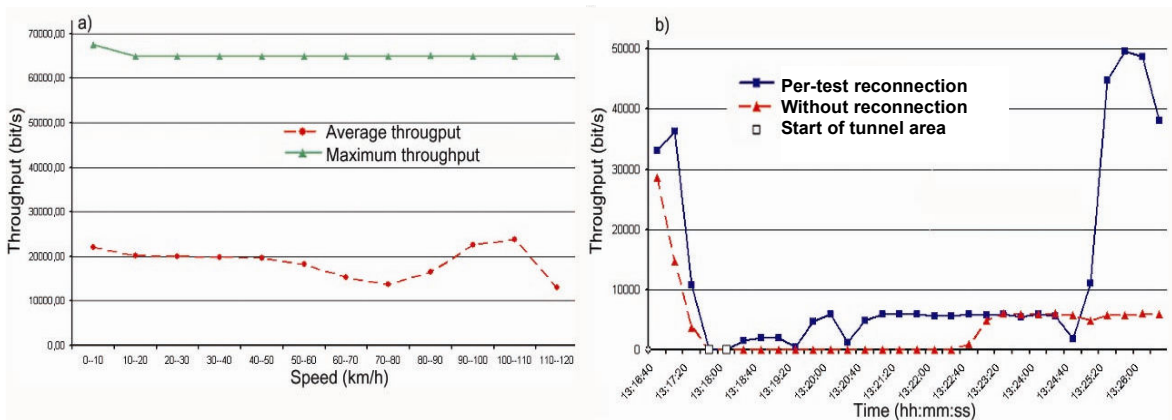


Figure 5. a) Throughput as a function of speed and b) impact of tunnels on bandwidth

As it can be observed in Fig. 5b, the best connection management strategy is to reconnect when traffic quality drops beyond a certain level (around 64 kbit/s).

- **Great dependability on the 3G network:** throughput fluctuates considerably (Fig. 6d) and this holds for all services (Fig. 6 a to c). It drops almost to zero when entering certain areas of the city of Madrid. Average network availability (Table 2) is only 58,28 %, although proper connection management at the car gateway (i.e., per-service reconnection) improves this availability by around 10 to 25% (Table 2).
- **Perceived quality:** video quality was measured qualitatively using a video player and quantitatively by minimum throughput threshold. Fluctuations below 320 kbit/s resulted in loss of audio and image freezings.

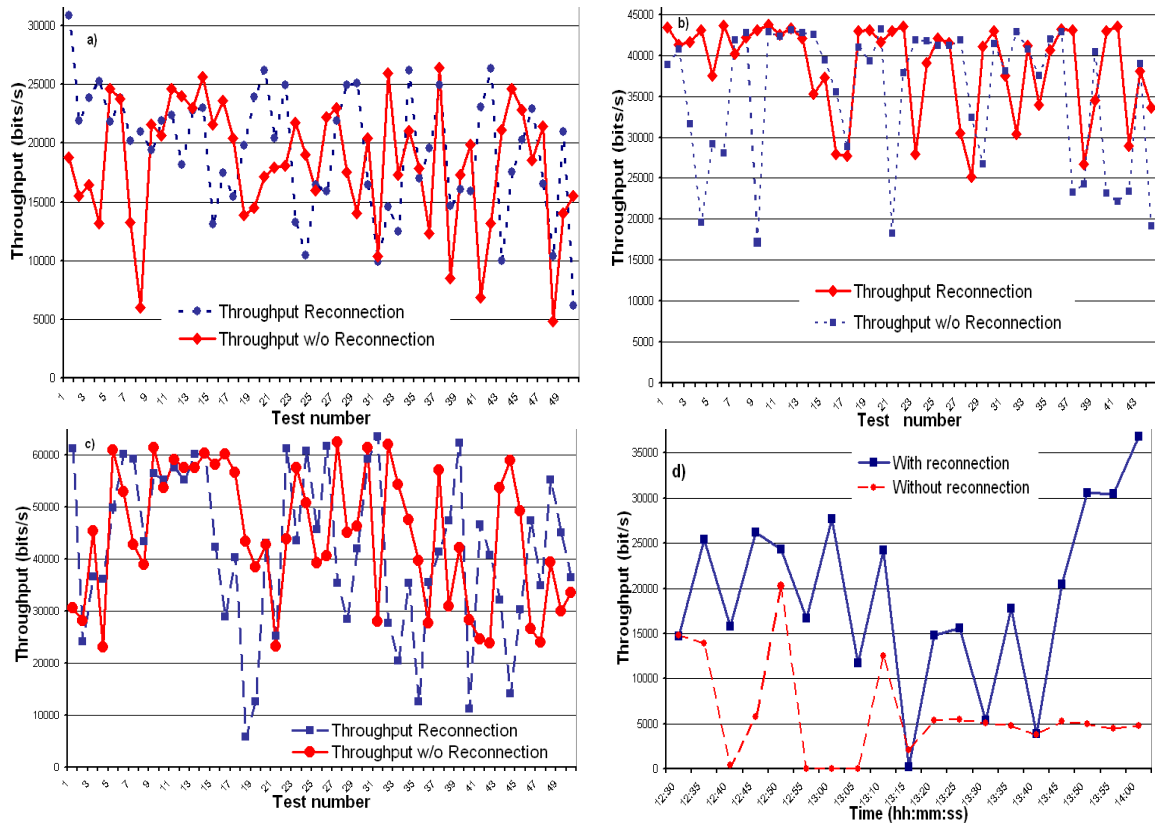


Figure 6. Throughput for a) web, b) video streaming, c) file download and b) 90-minute test

C. ALWAYS-ON CONSIDERATIONS

In the ABC scenario described in Section III.B, a service is provided from/to a vehicle through one or several wireless and mobile communication technologies transparently to both the service and to the user. In the DRIVE testbed, these technologies are currently WiFi (IEEE 802.11b/g) and 3G (HSDPA). Much research effort is being put into Vehicular Ad-Hoc Networks (VANET), which are based on short-range communications technologies. When thinking of services, it is not realistic to envision VANET communication for bandwidth-critical or greedy applications, because of two main facts: the bandwidth allocated for the future vehicular WiFi standard and the contact time. Currently, the Institute of Electrical and Electronics Engineers (IEEE) is standardizing an amendment to the WiFi family (802.11) for vehicles. This amendment – 802.11p – is also

known as Wireless Access in Vehicular Environments (WAVE). The spectrum of WAVE allocates 10 MHz channels for safety (1 channel), service (4), control (1) and long range/high power (1). As for the contact time, assuming that the WiFi range is 300 m, when driving at 75 km/h the contact period is of 30 s. With the overhead of association, DHCP and Internet connection, this period is even shorter than 30 s. Moreover, signal strength is a function of distance, which has an impact on the packet error rate as well as on the effective throughput of the connection. However, ABC will allow to switch between a VANET and an infrastructure network (e.g., 3G cellular) when a QoS metric exceeds a given threshold.

V. CONCLUSIONS AND FUTURE WORK

It seems clear that future vehicular services will require QoS. We have measured the network-QoS provided by 3G to critical applications in a simple vehicular scenario, and have obtained poor performance. The QoS measured and perceived in vehicles depend dramatically on the characteristics and availability of the 3G network. Such network dependability is a strong factor towards the integration of ABC in future on-board communications nodes. Future work encompasses the experimental measurement of QoS for vehicular services with dual WiFi/3G on-board communications node according to the ABC philosophy, which is expected to increase QoS even in severe network conditions.

ACKNOWLEDGMENT

This work was supported by Telefónica I+D and by the Spanish Government (MARTA project). The authors wish to thank E. Barea and C. Fuentes for their support in testing.

REFERENCES

- [1] ITU-T Recommendation E.800, "Telephone network and ISDN quality of service, network management and traffic engineering", 1994.
- [2] C. Pinart *et al.*, "DRIVE: a reconfigurable testbed for advanced vehicular services and communications", 1st WEEDEV Workshop, in Proc. TRIDENTCOM 2008, March 2008.
- [3] IETF RFC 3945, "Generalized Multi-Protocol Label Switching (GMPLS) architecture", October 2004.
- [4] IETF RFC 2998, "A framework for integrated services operation over DiffServ networks", Nov. 2000.
- [5] Ch. Makaya, S. Pierre, "An interworking architecture for heterogeneous IP wireless networks", in Proc. 3rd Int'l Conf. on Wireless and Mobile Communications, Mar. 2007.
- [6] D. J. Wright, "Maintaining QoS during handover among multiple wireless access technologies", in Proc. Int'l Conf. on the Management of Mobile Business, July 2007.
- [7] E. Marilly *et al.*, "Service level agreements: a main challenge for next generation networks", in Proc. 2nd Europ. Conf. on Universal Multiservice Networks, April 2002.
- [8] V. Sander *et al.*, "Networking Issues for Grid Infrastructure", Global Grid Forum GFD.37 (Informational), November 2004.
- [9] P. Gill, M. Arlitt, Z. Li, A. Mahanti, "YouTube traffic characterization: a view from the edge", in Proc. 7th ACM SIGCOMM conference on Internet measurement, Oct. 2007.
- [10] J. Derksen, R. Jansen, M. Maijala, E. Westerberg, "HSDPA performance and evolution", in Ericsson Review, vol. 83, no. 3, pp. 117-120, 2006.