Content Delivery Improvement by Satellite

Helmut Hlavacs

Institute for Computer Science and Business Informatics
University of Vienna
Lenaug. 2/8, Vienna, Austria
helmut.hlavacs@univie.ac.at

ABSTRACT

In this paper the CODIS content delivery network is described and its QoS is evaluated. CODIS is based on a high-speed satellite backbone which bypasses the Internet and interconnects four sites in Europe. Due to the inherent broadcast mechanism of satellites, CODIS efficiently distributes large amounts of data to an arbitrary number of receivers situated in its footprint over a single hop. The distributed data consists mainly of multimedia content like Web pages or IP-based video streams, which is replicated at different sites as close to the end-users as possible, thus increasing the end-user-perceived QoS. As users, both professional users like broadcasters, ISPs or multimedia companies, as well as consumers sitting at home are envisioned.

KEY WORDS

Satellite, Content Delivery Network, Quality of Service, end-user-perceived QoS

1 Introduction

The next generation Internet will carry a mixture of pure data, Web services and multimedia real-time traffic like Voice over IP, live broadcasting or video on demand. Unfortunately the current IPv4 Internet protocol suite works in best effort mode only and cannot guarantee timely delivery for real-time traffic, thus imposing the danger that this type of traffic may be severely disturbed by overloaded routers. Things become worse if the transported real-time presentation needs high bitrates, as is necessary for instance, for TV resolution, requiring bitrates of more than 1.5 Mbit/s when using modern encoders like H-264 [3]. Proposed improvements like RSVP [1] or IPv6 [2] are only supported by a small subset of the existing Internet routers yet and thus offer no improvement.

In fact the farther away the multimedia source is from the receiver, the higher the probability gets that real-time traffic is affected by cross-traffic. *Content delivery networks* (CDN) are a means for providing a better network QoS to paying customers, such that for instance multimedia streaming and Web access are affected by the ever increasing Internet cross-traffic no more. This can be achieved on the one hand by leasing private network links, thus excluding the traffic from the rest of the Internet, and/or by reducing the number of hops between sender and receiver. This can be done using a mixture of bypassing frequent Internet routes with private links and replicating multimedia content at locations being spread all over the continent or even the world. The latter scheme provides large caches being close to paying customers, for instance customers of an Internet service provider (ISP) or employees of large companies. The problem for distributing large amounts of data replicas to the caches remains. Using unicast wastes enormous amounts of bandwidth, while using multicast over the Internet, like in the case of RSVP and IPv6, still is supported by a small subset of Internet routers only, which is called MBone.[1]

In this paper the content delivery network CODIS[2] (Content Delivery Improvement by Satellite) is presented. CODIS uses the high-bandwidth capacity of GEO satellites as a single hop private network for content distribution. CODIS exploits the natural multicasting capabilities of such satellites for reaching arbitrary numbers of receiver stations situated in their footprints. Furthermore, CODIS combines IP multimedia data like streaming video or Web data with DVB-T/DVB-S based TV broadcast efficiently, which can be exploited for instance by future *multimedia home platform* (MHP) set-top boxes.[3]

2 The CODIS CDN

CODIS is an IST[4] (Information Society Technologies) project supported by the European Commission.[5] The CODIS project is lead by the satellite manufacturer Alcatel Space situated in Toulouse, France. Other partners are given by the research institute Télédiffusion de France (TDF) managing part the television infrastructure of France, and the Institut für Rundfunktechnik (IRT) carrying out research for the public broadcasting companies of Germany, Austria and Switzerland, both TDF and IRT running DVB-T (digital television being broadcast from terrestrial infrastructure) test stations. Additional partners are the measurement equipment manufacturer Rohde&Schwarz who provided DVB-S (digital television being broadcast over satellite) test equipment, and the French start-up company ActiVia Networks, who provided the CDN software. The French space agency CNES provided the satellite equipment. The task of the University of Vienna was to define a framework for measuring the QoS of CODIS as perceived by its end-users, and to carry out and analyze the measurements. The CODIS end-users are projected to be professionals who run the CDN and consumers sitting at their homes or at work who actually consume the content.

The initial aim was to use the STENTOR satellite created by Alcatel Space and Astrium. STENTOR was an innovation in the sense that it provided an on-board DVB-S multiplexer being capable of receiving several DVB-S transponders and multiplexing them into one single DVB-S stream which is

[1] http://www.lbl.gov/web/MBONE.html

[2] http://www.codis-satellite.com/

[3] http://www.mhp.org/

[4] http://www.cordis.lu/ist/

[5] http://www.cordis.lu/

then sent back to earth again. Unfortunately STENTOR was destroyed during launch due to the failure of the Ariane V in December 2002. As a fallback solution the consortium used two other commercial satellites already being available, namely Telecom 2D and Atlantic Bird 2, both being situated at 8 degree West.

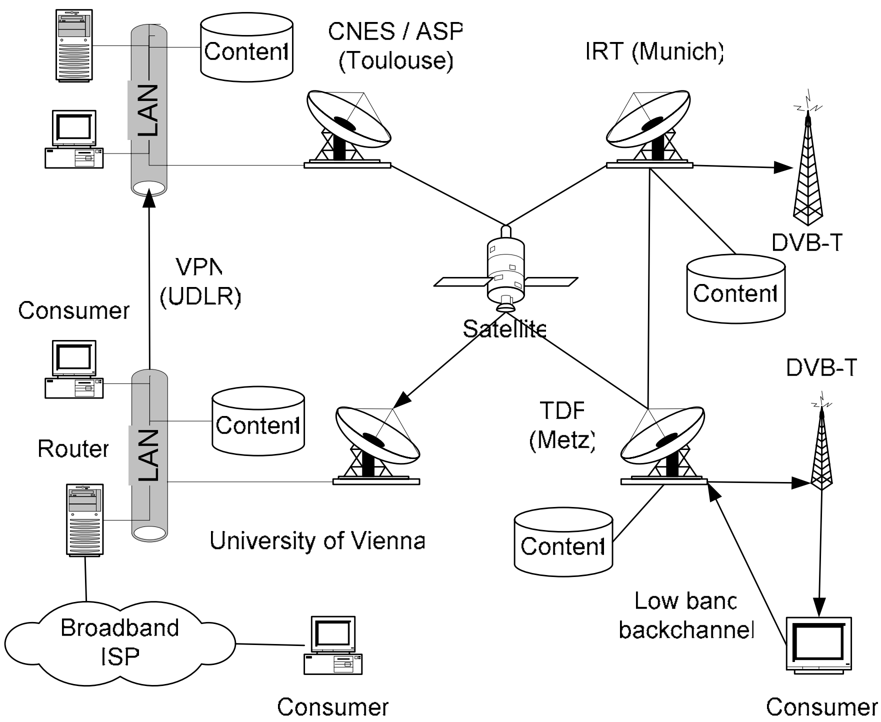

Figure 1: The CODIS mesh.

The CODIS structure is shown in Figure 1. At each site a local content cache is installed. Also, each site is able to send content to or receive content from other sites via the used satellite, with the exception of the Vienna QoS measurement site (see Section 5.2), which only receives data from the satellite, but sends data via a virtual private network (VPN) based on the UDLR protocol to Toulouse. The architectures of Vienna and Toulouse resemble an ISP-like infrastructure, while Munich and Metz represent the case of broadcasters using DVB-T for delivering digital TV enriched with IP-based contents to their customers.

2.1 Internet Service Providers

One aspect of CODIS represents ISP core networks. In this scenario, ISPs receive content which can be distributed to their paying customers. No DVB services are included here. However, as such a high-speed network is completely managed by only one entity, and the number of hops between end-users and the ISP servers is likely to be small, IP based pay-per-view services can easily be implemented by the ISP, offering for example video on demand in TV resolution.

The Alcatel platform has been designed to resemble such a network. Here, the core network is attached to a central IP router providing unicast and multicast services. The router interconnects the local servers to the customers on the one hand, and the satellite backbone on the other. Content is either stored on the local Web and video server containing original content, or on the local cache which is called *Generic Server*, a technology which has been developed by the CODIS consortium. Furthermore, a local DNS server controlling the local domain *codis-satellite.net* provides routing services.

IP traffic is sent to the satellite using an IP/MPEG-2 encapsulator, and received via an IPricot IP receiver. Both are connected to the satellite out-door station which has been provided by CNES.

2.2 Broadcasters

CODIS is also designed to cooperate with future digital television broadcasting being based on the DVB-T standard [5]. In this scenario, broadcasting companies may include IP based data into the broadcast DVB signal. This may include short trailers that can be viewed on demand, Web data, games, interaction with live programmes like voting, or teletext like information for the MHP set-top box. Prerecorded videos may be stored in the local Generic Server of the broadcasters and may be accessed by the users. This can be done, for instance, by inserting data into a data carousel that is constantly broadcast to the MHP set-top boxes. Depending on the set-top box capabilities, the presentations may then be viewed as they are currently broadcast, or may be stored on the local MHP hard disk for on demand access later. It is worth noting that in this scenario, the focus is put more on broadcasting predefined audio-visual content to large audiences than on on-demand content access by individuals. The latter requires a back-channel from the user to the broadcaster, which may be realized using modems, ADSL, GPRS etc.

CODIS supports this scenario as IP based content can easily be exchanged between broadcasting companies, and between different local branches of national broadcasters by using the CODIS CDN infrastructure.

In our experiments Munich and Metz represented the broadcaster side. Metz used its DVB-T infrastructure consisting of the main transmitters in Luttange and Scy-Chazelles, a repeater site in St Julien and two gap-fillers inside the TDF building. The connection to the satellite was established using standard IP/DVB gateways, the Metz site is capable of distributing up to 20 Mbit/s of IP data to end-users.

3 Content Delivery

Content delivery to the consumer requires two steps. At first the content is distributed from the site of origin to the targeted caches. Here, the CDN operator chooses the content to be published, and then multicasts it to the receiver caches. This operation is called *push* or *publication*. On the other hand, content may only be replicated as a result of a consumer request. In this scheme, the content is at least stored at the cache nearest to the requesting user, resulting in a *pull* operation.

Once a user requests certain content, the CDN system must find the cache closest to the user which contains the content. Here, the CDN re-interprets the user supplied URL and may also initiate content transfers. As a result, the content may be served from a server being different from the one that was initially specified in the URL. However, this must be done in a completely transparent manner, without users noticing the CDN layer.

Content distribution management is carried out using the Constellation Manager software. This software is Web based and provides means for the CDN operator managing the CDN and the CDN customers, as well as means for content providers to push their content to other sites. Also the Constellation Manager allows the real-time monitoring of the CDN state.

3.1 Content Distribution

The CODIS CDN focuses on two types of content: Web pages or subtrees of Web sites, and video encoded with the MPEG-4 visual layer codec. These two types require two different operation sequences: For pushing Web page subtrees the following operations are necessary:

1. On the Generic Server the program wget retrieves the required content from the origin Web server.
2. Then the files are put into an archive using tar, and the archive is compressed with gzip.
3. Then the actual multicast is initiated using SAT-RMTP [7].
4. At the receiving Generic Server, the archive is decompressed and untared.
5. Then the Web pages are put into the squid proxy server.

Points 1, 2, 4, and 5 comprise the CDN overhead, while only point 3 actually uses the satellite (network) connection. It must be pointed out that HTML Web pages and other text files enable high compression factors, while embedded graphics files like JPEG or GIF already are compressed and thus will not be reduced in size when being compressed again. For MPEG-4 videos, the push operation is much simpler:

1. Move the MPEG-4 file from the origin server to the Generic Server.
2. On the Generic Server compress the file with gzip.
3. Carry out the transfer over the network.
4. On the receiving Generic Server uncompress the file and put it into the streaming server's movie directory.

Here, the CDN overhead is given by points 1, 2 and 4. However, as MPEG-4 already is highly compressed, the achieved compression rates will be quite low, and the actually data transfer will be a significant part of the overall push time.

3.2 Content Routing

Content routing is the task of automatically and transparently redirecting content accesses to replicas of requested content stored in caches being closest the requesting user. For this purpose, CODIS manages its own DNS domain *codis-satellite.net*, using an authoritative server in Toulouse, and slave servers in Metz, Munich and Vienna. When trying to resolve a DNS name of this domain, the CODIS DNS server replaces the user supplied DNS name with the DNS name of an appropriate Generic Server closest to the user, i.e., Metz for Metz users etc. This procedure is called *CDNization*.

4 Security

As CODIS transports proprietary multimedia content with restricted access rights, making sure that only authorized people have access to the transported content is of utmost importance. For this purpose CODIS uses a two-step scheme. The first security layer is given by the standard DVB-S Conditional Access (CA) system. This is a DVB standard based on PCMCIA, where special PCMCIA cards using the so-called Common Interface can be used for decrypting content encrypted by the CA system. However, this grants access to all content sent over the satellite which is encrypted using the same CA module, and CODIS should be able to multicast its content to arbitrary subgroups of such users.

Thus the second security layer is given by an extension of IPSec called SatIPSec [4]. SatIPSec for CODIS ensures a secure channel between each *receiving* Generic Server (here acting as clients) with each other Generic Server (which may act as the multicast originating server). This is done in three steps. First each client realizes mutual authentication with the originating server, establishing also a secure channel with the server. In the second phase the server sends a table called Security Association (SA) containing security related parameters to the client by using unicast. Additionally, if the client is member of specific multicasting groups, i.e. news, sport etc., the server sends a Control SA to the client, one for each group the client is member of. Thirdly, if the server wants to do a publication for a specific group of clients, it sends a Data-SA to all authorized clients owning the Control SA for the particular service, the clients later use this Data-SA for decrypting the respective multicast data.

5 Quality of Service

The task of the University of Vienna was to define a framework for measuring the QoS of CODIS as the CODIS end-users would see it, as well as to carry out these measurements (within well defined experiments) and analyze their results. For this purpose we have set up a QoS measurement platform and either used ready to use measurement tools, or on the other hand programmed our own tools due to various limitations of existing tools.

5.1 The QoS Measurement Framework

For creating a framework [8] for measuring the QoS of the CODIS CDN, we put our focus not only on standard network metrics, but also the user-perceived QoS. Thus, CODIS must satisfy the needs for professional users like broadcasters, ISPs, etc. but also the ones of consumers sitting at home.

Figure 2 shows the dependence of the user-perceived QoS on the network QoS, the latter being the sum of the QoS of the individual protocol layers.

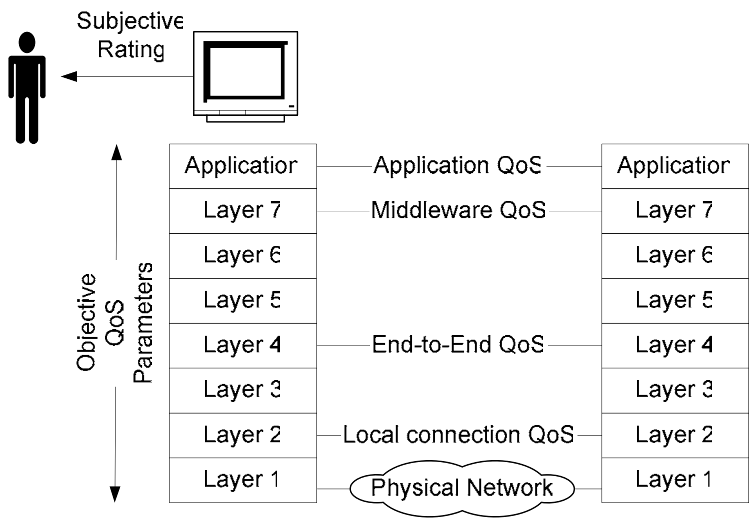

Figure 2: Connection between end-user-perceived QoS and network QoS.

As subjective experiments usually require an enormous amount of time and a large number of test subjects, we first tried to find as many mappings from objective to subjective QoS as possible. Figure 3 shows such an example taken from [10]. In their experiments, the authors asked 30 subjects to subjectively rate the perceived waiting time when downloading Web pages. The subjects chose a rating from the set of possible ratings *high*, *average* and *low*. It is interesting that the *low* rating category agrees with numerous results found in the scientific literature about waiting time rating. The important limits found there, are for instance, 4 seconds, 6 seconds, 8 seconds and 11 seconds. At these limits, the other results as well as [10] report a significant lowering of the perceived QoS, the absolute maximum allowed being 11 seconds.

Figure 4 shows the network layers we chose for carrying out QoS measurements. This decision depended on several factors. First, they should represent important CODIS aspects, like pure IP transfer, the CODIS CDN part and the application part (Web, MPEG-4). Second, the chosen layers in sum should be responsible for the user-perceived QoS, thus include application, transport or network and MAC layers. Third, we should be able to carry out measurements at these levels without being forced to change network drivers or the TCP/IP stack of the used operating systems. Instead, we relied mostly on off-the-shelf tools being either pure software, or hardware provided by our CODIS partners. However, since the existing tools where not capable to measure all the metrics we aimed for, we chose to create a set of CODIS measurement tools ourselves, like CODIS RTSP for measuring MPEG-4 streaming, and CODIS Net for measuring TCP/UDP metrics.

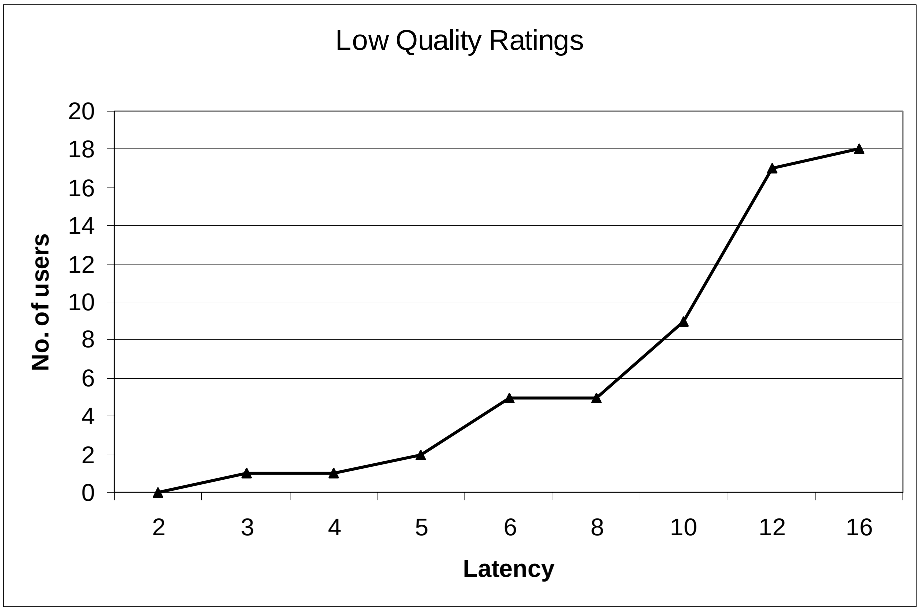

Figure 3: Mapping of experienced waiting time (latency) to *low* subjective QoS.

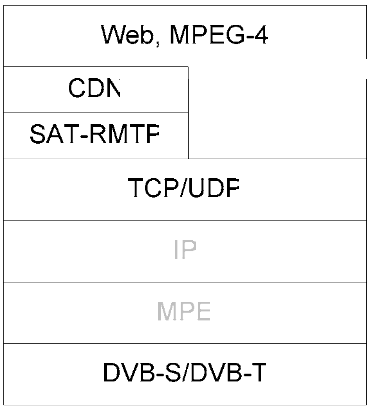

Figure 4: CODIS network stack: Middleware and applications using IP over DVB-S.

5.2 Vienna QoS Measurement Platform

Additionally to defining the QoS metrics we set up a platform for measuring the CODIS QoS seen in Figure 5, which has been implemented in a rack placed in a cellar below the satellite dish.

The main DVB-S signal is received by the satellite dish, which is connected to four pieces of hardware. The IPricot receiver pulls the IP packets out of this signal and passes them on to the workgroup switch. The DVB-S signal is also forwarded from the IPricot to the DVB-S demodulator Newtec, which results in a pure MPEG-2 data stream from there on transported over an ASI interface. The pure MPEG-2 stream then is led into the analyzers DVMD and DVQ, both provided by Rohde&Schwarz. The DVMD analyses the QoS of the pure MPEG-2 QoS, flagging events and errors as defined in [6], the DVQ analyses the picture quality of the videos embedded into the MPEG-2 stream. The DVB-S signals are also provided to a Fujitsu/Siemens Activity set-top box, and a Windows XP based PC containing a DVB-S receiver card.

In the middle the Windows based PC platform is shown where Windows related measurement software has been used. A Linux based firewall/router connects the private network with the Internet, which made possible the remote control of most experiments from our institute building, but which also served as the CDN Generic Server and which has been used for Linux based QoS experiments.

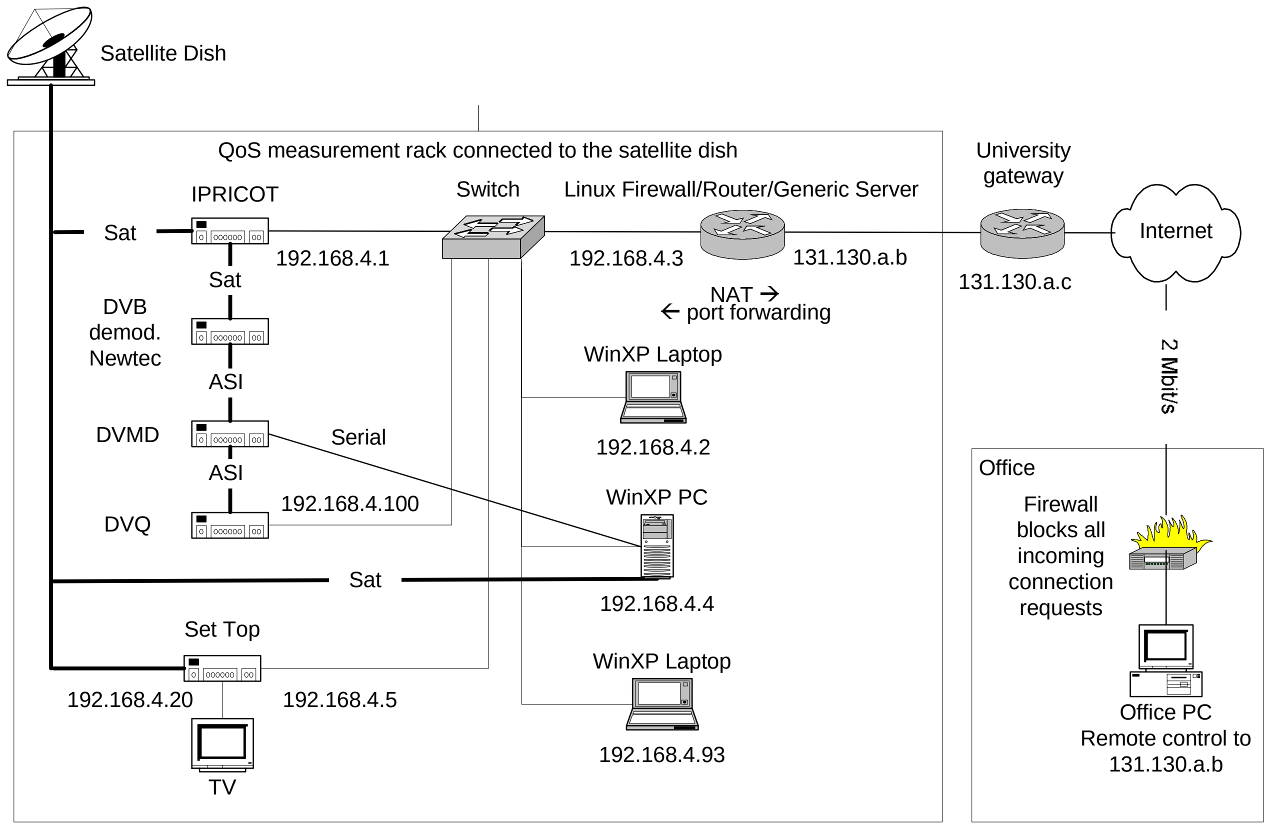

Figure 5: Vienna QoS measurement platform.

5.3 Measurement Results

In addition to planning experiments which included the satellite connection, we also carried out numerous experiments which included other CODIS parts like LAN, MAN or DVB-T. Also, we carried out experiments over the Internet in order to be able to compare the CODIS QoS to the one provided by its common alternative. In total, we carried our experiments at the CODIS main sites in Vienna, Toulouse, Munich and Metz, as well as at the University of Mallorca, at an ISP provider in Rome and in Amsterdam (WAN experiments).

5.3.1 DVB-S

For our QoS experiments [9], DVB-S streams were broadcast from the IRT station and received by the equipment located in Vienna. There, the incoming stream was fed into the Rohde&Schwarz

measurement devices, the DVMD and the DVQ. The software Stream Explorer was used to analyze the stream and to log the observed errors to files, which were evaluated later. There are three error priorities defined for DVB-S [6]. Errors of priority 1 result in an immediate loss of sound and picture. If such errors occur, it is not necessary to look for errors at layers 2 and 3. Errors at priority 2 may result in intact programmes, but important meta-information like information tables are not received or received but not in time, and will result in disturbed services. Errors at priority 3 are application specific and will result in disturbance of only specific applications.

Figure 6 shows a typical result for errors at priority 2. Generally there are no errors, except that the sender does not include the Network Information Table (NIT) into the MPEG-2 stream, resulting in permanent NIT-error messages. The conclusion here is that the satellite connection is very stable and yields a reliable multi-purpose (TV and IP data) high-bandwidth network. However, during a heavy thunderstorm in Vienna, we once observed a loss of connection for about 10 minutes

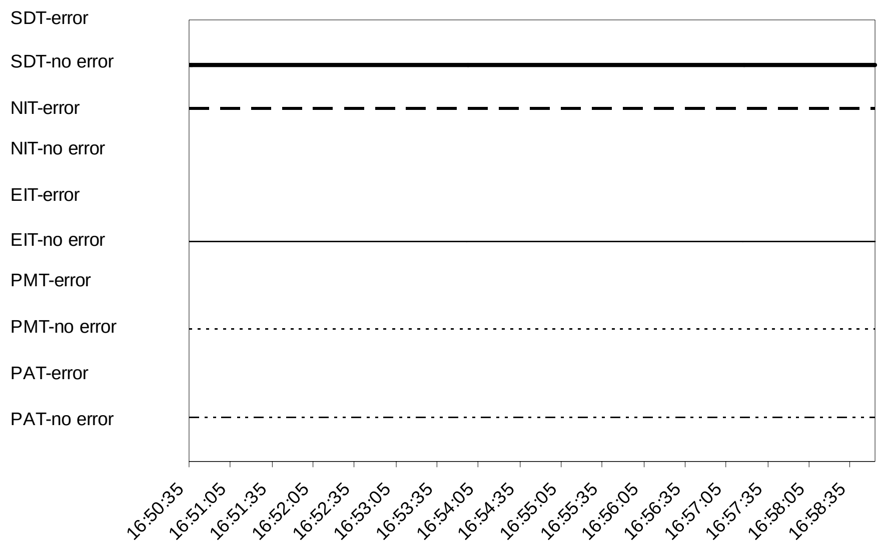

Figure 6: Typical priority 2 errors.

5.3.2 IP

At the IP/TCP/UDP level we measured several different metrics for the DVB-S and DVB-T based CODIS networks, but also for different network types over the Internet like LAN (FastEthernet), MAN (inside Vienna), and WAN (between Vienna and Toulouse, Metz, Rome and Mallorca).

Table 1 shows the number of hops a packet must pass by between sender and receiver. It can be seen that, although being a continent wide network, the DVB-S has only 3 hops, while terrestrial WANs need up to 16 hops over different routers.

We have measured the round trip time (RTT) for various network types. In Figure 7 the empirical cumulative distribution functions (ECDF) for them are shown. For measuring the RTT we used the programs ping and the CODIS Net tool. The RTT for the LANs could not be determined accurately, as the timer resolutions for both ping and CODIS Net are 10^{-3} s, and thus this is the upper bound which is stated. It is interesting to see that some Internet WAN connections show a high variance (Toulouse, Rome), while the others are quite stable. One particularity of the used DVB-T equipment can be seen. The RTT for a moderately loaded DVB-T network is much lower than the one of a lightly loaded network. This phenomenon is due to the used DVB-T multiplexer which schedules IP packets more advantageous in case a minimum data rate is observed (optimized for continuous video streams transfer, not bursty IP traffic).

Table 1: Network hop counts.

ID	Network	From	To	Hop Count
1	CDN LAN	Vienna LAN	Vienna LAN	1
2	CDN MAN	Vienna CDN	Vienna Home	9
3	DVB-S	Toulouse	Vienna	3
4	DVB-T	Metz CDN	Metz Home	3
5	Internet	Vienna	Toulouse	16
6	Internet	Vienna	Rome	15
7	Internet	Vienna	Mallorca	>12

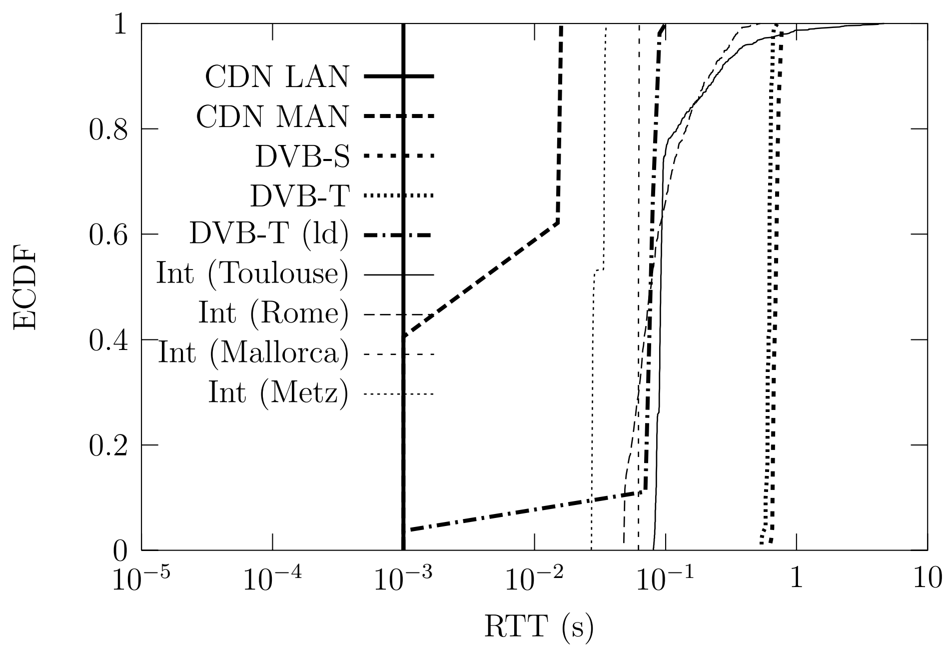

Figure 7: End-to-end round trip time for various network types.

The general packet loss rate for the DVB networks has been measured with CODIS Net and is shown in Figure 8. Generally, the loss rate rises quickly as soon as the bottleneck bandwidth is approached. An exception is given by the case DVB-S (problem), where only a 10 Mbit/s hub was used between the IPricot received and the measurement PC, and additionally a high bitrate multicast was taking away a considerable share of the bandwidth. For DVB-T we used two different predefined maximum bitrates: 2.5 Mbit/s for DVB-T (1) and 1.5 Mbit/s for DVB-T (2).

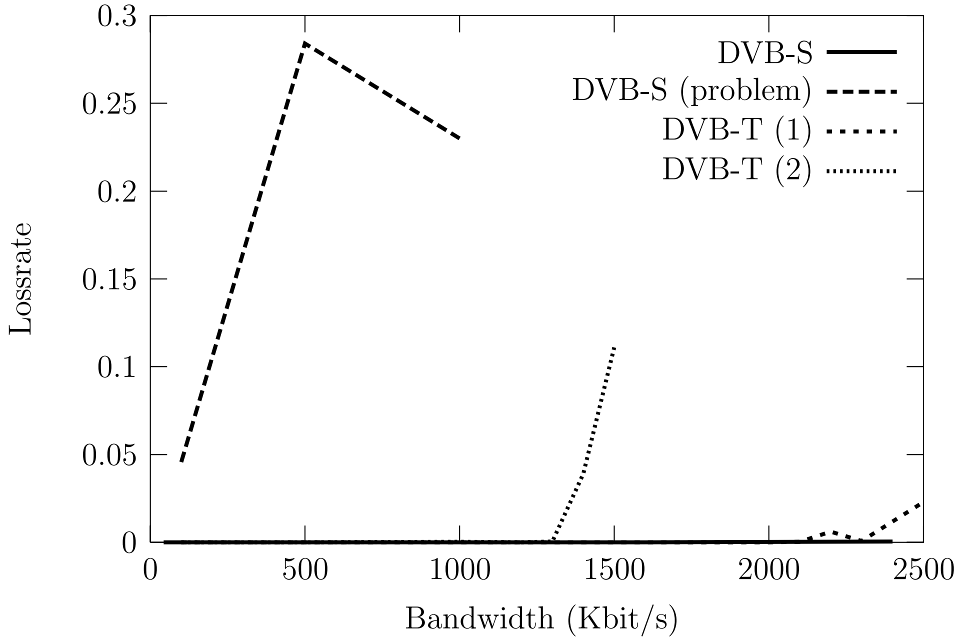

Figure 8: Packet loss rates for the main CODIS networks.

5.3.3 CDN

For CODIS we have carried out download experiments and clocked the publication times for typical Web and MPEG-4 content. The result can be seen in Figure 9. The bold lines show the download times for pushing content of different type and size from two sources, IRT and TDF. The thin lines denote the pure transmission times of the compressed data. OD denotes MPEG-4 content, Web denotes Web subtrees containing text files and bitmap data, both can be highly compressed. IRT OD (ld) means that in parallel to the pure IP transmission, IRT was sending an MPEG-2 DVB-S live broadcast over their carrier. IRT used a 4.8 Mbit/s carried, while TDF used a 2 Mbit/s carrier, but restricted themselves to 1 Mbit/s CDN sending bitrate.

In another experiment we observed the utilization of standard PCs when serving a large number of IP based real-time video streams (MPEG-4 over RTP). As was expected, the most severe bottleneck was given by the serving hard disk (Figure 10), while the memory and CPU remained only lightly loaded. The figure shows the hard disk utilization when streaming different files (thick lines) and always the same file (thin lines) for videos encoded with different bitrates. The tough case is streaming different files, while streaming always the same file does not increase the utilization significantly.

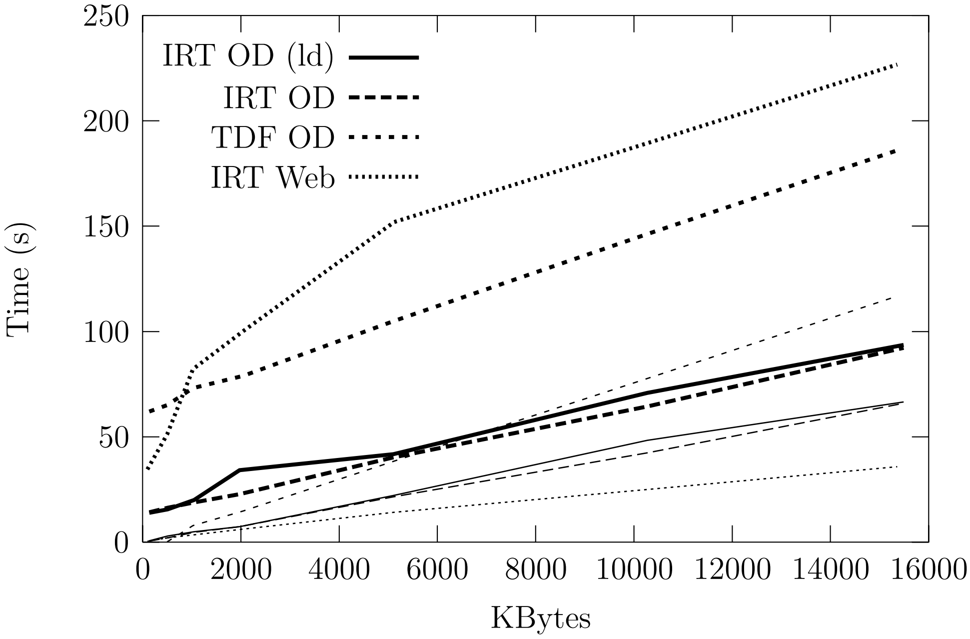

Figure 9: CDN overall publication and pure data transmission times.

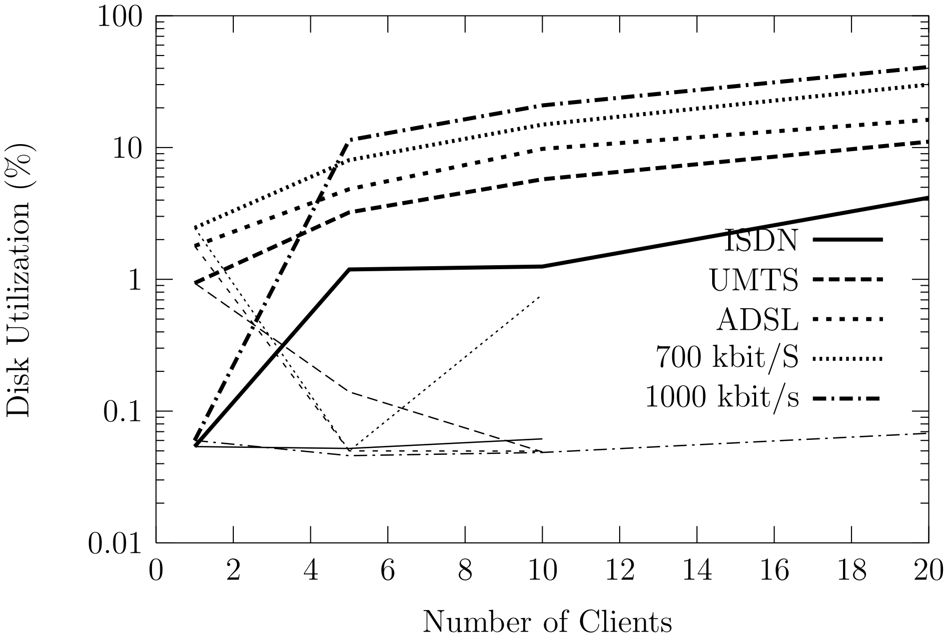

Figure 10: Disk usage when serving different target networks.

5.3.4 User Access

Figure 11 shows the response time of the ASP e-learning application, which is a Web site from Alcatel Space, additionally containing graphics, embedded streaming videos and MS Power Point presentations. In the diagram, (Cache) means that the browser cache was enabled, thus decreasing the download-time of re-referenced Web pages accordingly. In the (No Cache) case, the browser cache has been disabled. Results are stated for downloading over LAN, a mixture of Internet (upstream) and satellite (down stream) denoted as (Int./DVB-S), and over a Viennese ISP (CDN MAN), offering Internet access via cable.

It can be seen that the used Web pages are mostly loaded within the high quality limit, the Web application can thus be used using all investigated network types. Of course, the cached LAN cannot be beaten and responds without visible delay. But also the Internet/satellite case results in satisfactory performance.

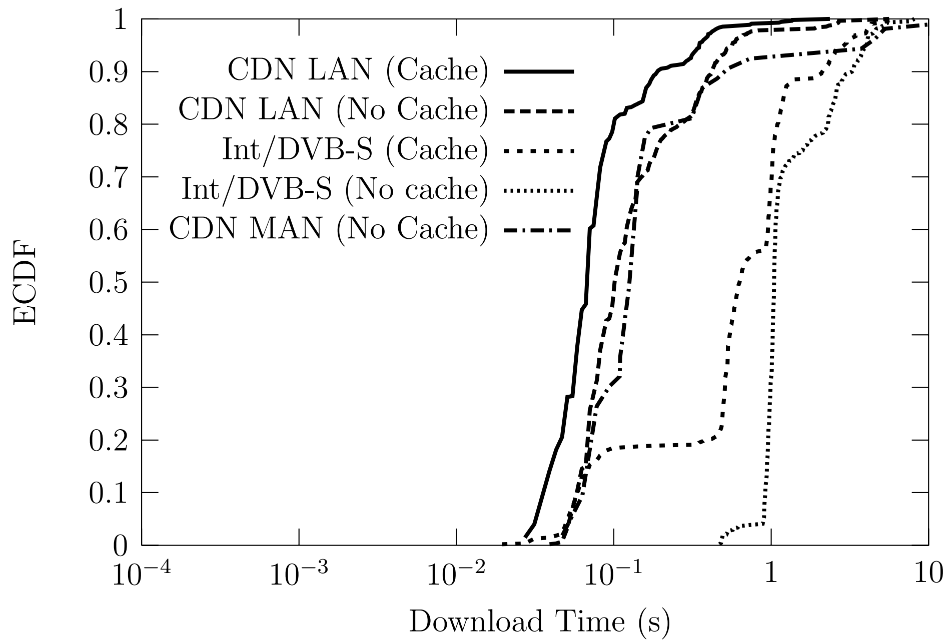

Figure 11: Response time for ASP e-learning application.

For streaming MPEG-4 we use the standardized protocols RTSP for video control and RTP/RTCP as underlying real-time streaming protocols. The meaning of RTSP is the one of a remote control, issuing start and stop requests over TCP. The audio/video data itself is packed into UDP/RTP packets, and RTCP is used for exchanging sender and receiver reports (statistics) with each other.

Our tool CODIS RTSP can be used to automatically measure the end-user response time, which we have done for the network types shown in Figure 12. As streaming server, the Real Helix server was used, with exception of the CDN LAN case, where the Apple Darwin server was used. The network DVB-S/DVB-T uses the satellite link and additionally a DVB-T link for communication. The (ld) case again denotes a medium loaded DVB-T network which (as shown previously) exhibits lower latencies. It can be seen that problems may arise in the DVB-S/DVB-T case where the measured response times exceed the upper limit in roughly half the cases

6 Conclusion

In this paper the IST project CODIS is presented. CODIS is a content delivery network which uses a central satellite as a single-hop high-speed backbone for content distribution. The CODIS consortium focuses on professional users like broadcasters, ISPs or companies, as well as on end-users sitting at home and consuming the content using, for instance a DVB-T MHP. CODIS is able to provide IP based multimedia content like Web pages or MPEG-4 videos, in order to serve ISP users as well as to broadcast IP enriched DVB programmes to large audiences.

The consortium has developed, set up and run the CDN. In numerous experiments, the CODIS QoS has been evaluated, measuring not only standard QoS metrics, but also trying to interpret the QoS from the view of the CODIS end-users. The results show that the satellite based distribution is well suited for moving for the CDN tasks.

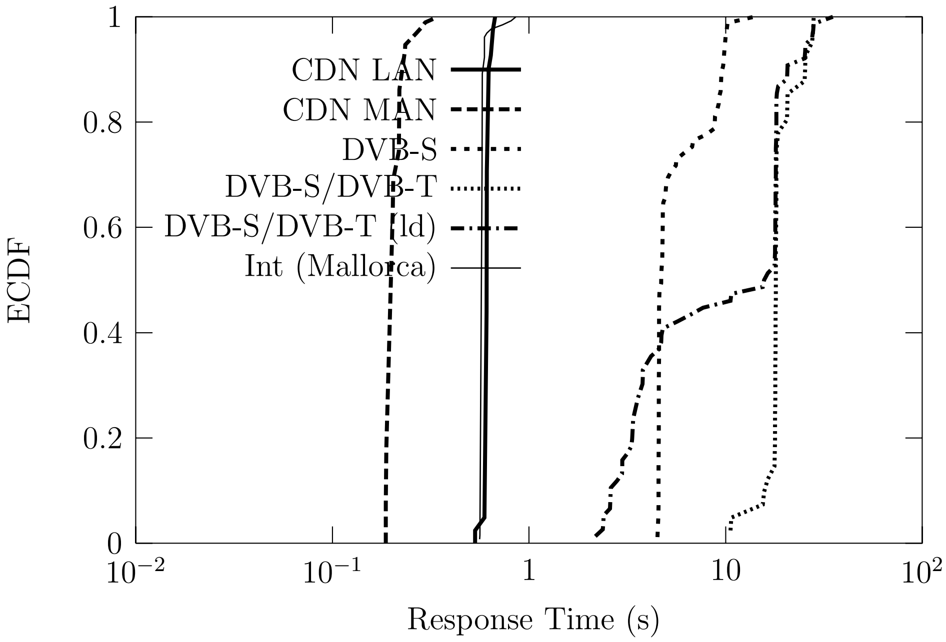

Figure 12. MPEG-4 response time measured by CODIS RTSP.

7 References

[1] Braden et al., *Resource Reservation Protocol (RSVP)*, RFC 2205, IETF, 1997.

[2] S. Deering, R. Hinden, *Internet Protocol, Version 6 (IPv6)*, RFC 2460, IETF, 1998.

[3] *Advanced Video Coding*, ISO/IEC 14496-10 (H-264), 2003.

[4] S. Josset, L. Duquerroy, M. Önen, M. Annoni, G. Boiero, N. Salis, *Satellite IPSec: an optimized way of securing multicast satellite communications*, in: ISSE 2002, Paris, France, Oct. 2002.

[5] U. Reimers, *Digital Video Broadcasting*, Springer Berlin Heidelberg, 2000.

[6] Digital Video Broadcasting (DVB): *Measurement Guidelines for DVB Systems.* ETSI TR 101 290 V1.2.1 (2001-05), ETSI, 2001.

[7] M. Koyabe , G. Fairhurst, A. Matthews, *Satellite Reliable Multicast Transport Protocol (SAT-RMTP): Specification and Functionality*. Technical Report: GEOC-UOA-2400-2,

[8] H. Hlavacs et al., IST 2000-30066 CODIS Deliverable D2300: *QoS Analysis Method*, http://www.codis-satellite.com/deliverables.php.

[9] H. Hlavacs et al., IST 2000-30066 CODIS Deliverable D3400: *QoS Measurements*, http://www.codis-satellite.com/deliverables.php.

[10] A. Bouch, A. Kuchinsky, N. Bhatti, *Quality is in the Eye of the Beholder: Meeting User's Requirements for Internet Quality of Service*, Technical Report HPL-2000-4, HP Laboratories Palo Alto, January 2000.