Information and Communication Technologies (ICT) Programme

Project N°: FP7-ICT- 214063

SEA

Deliverable D4.4.1

sNMG system prototype description

Author(s): M. Naumann, K. T. Pijarn (Nomor Research GmbH), C. Kontopoulos (Vodafone), K. Grüneberg (HHI), Th. Zahariadis (Synelixis), G/ Olmo, S. Zezza (POLITO)

Status -Version: Final

Date: 1 June 2009

Distribution - Confidentiality: Public

Code: SEA_D4.4.1_NOM_FF_20090601.doc

Abstract:

This document gives the overall specifications of the SEA Network Media Gateway (sNMG), which will be used for SEA demonstrations. The document analyses the requirements of the sNMG in terms of computational complexity, memory and storage requirements and describes the software solution for the implementation of the prototype to be presented as an output of the project.

It has to be mentioned that emphasis has been put on the system/prototype functionality and not on scaling issues, which are mainly related to deployment parameters.

Disclaimer

The SEA Consortium consists of the following companies:

No	Participant name	Participant short name	Country	Country
1	ST Microelectronics	STM	Co-ordinator	Italy
2	Synelixis Solutions	Synelixis	Contractor	Greece
3	Thomson Grass Valley	Thomson	Contractor	France
4	Philips Consumer Lifestyle	Philips	Contractor	Netherlands
5	Vodafone Panafon AEET	Vodafone	Contractor	Greece
6	Nomor Research	Nomor	Contractor	Germany
7	Fraunhofer HHI	HHI	Contractor	Germany
8	Politecnico di Torino	Polito	Contractor	Italy
9	Universidad Politécnica de Madrid	UPM	Contractor	Spain
10	University of California, Los Angeles	UCLA	Contractor	USA

This page has been intentionally left blank

Table of contents

Abbreviations

AQoS	*Adaptive Quality of Service*
AVC	*Advanced Video Coding*
MDC	*Multiple Description Coding*
AE	*Adaptation Engine*
AEM	*Adaptation Execution Module*
TAM	*Terminal Awareness Module*
ADM	*Adaptation Decision Module*
NAM	*Network Awareness Module*
NALU	*Network Abstraction Layer Unit*
RG	*Residential Gateway*
RS	*Redundant Slices*
SVC	*Scalable Video Coding*
MVC	*Multi Viewpoint Video Coding*
MDC	*Multi Description Coding*
UCD	*User Constraints Document*
UED	*User Environment Document*
RTSP	*Real Time Streaming Protocol*
P2P	*Peer to Peer*
sNMG	*Seamless Network Media Gateway*
sHMG	*Seamless Home Media Gateway*

1 Introduction

SEA aims to offer seamless video delivery, maintaining the integrity and wherever applicable, adapting and enriching the quality of the media across the whole distribution chain. SEA has proposed the abstract network architecture shown in Figure 1. We introduce two new types of Media-Aware Network Elements (MANE):

a) a seamless Home Media Gateway (sHMG), located at the edge of the extended home environment and

b) a seamless Network Media Gateway (sNMG) at the edge of the access networks, e.g. the 3GPP Service Architecture Evolution (SAE)

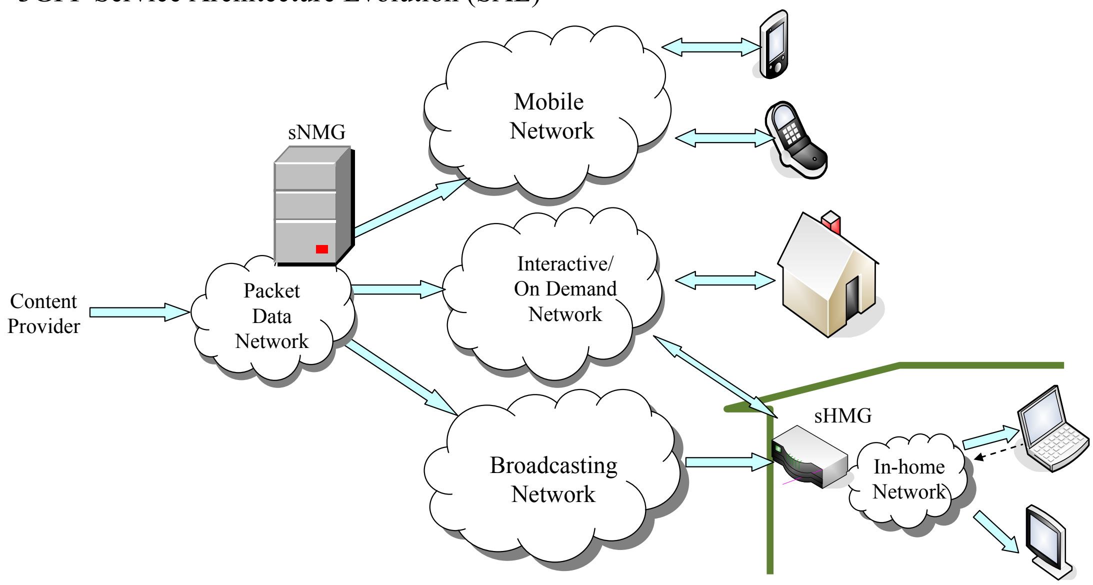

Figure 1: SEA abstract network architecture

The SEA content aware MANE edge nodes (sHMG and sNMG) are network-based components for SEA architecture and support the intelligent, seamless content distribution. They offer functions like network and terminal awareness, content enrichment and content protection. These MANE nodes can offer multimedia storage, dynamic content adaptation and enriched PQoS by dynamically combining multiple multimedia content layers from various sources. Moreover, as they have knowledge of the underlying networks, this information on the network conditions/ characteristics can be provided to and utilized by the Cross Layer Control mechanism and adapt the multimedia streams to the next network in the delivery path. For more details, the readers are referred to [1] and [2]

1.1 SEA mapping to the SAE architecture

A simplified overview of the 3GPP System Architecture Evolution (SAE) concept (see 3GPP TR 23.882 v1.6.1) together with the SEA's introduced MANE's is depicted in Figure 2. In a 3G/3.5G environment, the radio and the core networks are under the control of the mobile operator. Typically, if external content service providers want to offer their services not only on a best effort base, but on a well defined QoS, they are dependent on the SGi open interface to operator's service architecture. Thus, for the integration of SEA sNMG into SAE environment the reference points to the SAE anchor, the SGi interface is of major relevance. A clear definition and implementation of these access points will enable the assessment and easy integration of the considered SEA services into emerging service architectures.

For maintaining the QoS throughout the entire distribution path, it is also important to allow mapping of service QoS from the sNMG to the SAE's QoS policy. For this purpose, the SAE architecture allows pushing QoS profiles based on subscriber IDs into the network all the way to radio access networks. For this, an optional access to the Policy control and Charging Rules Function (PCRF) via the Rx interface is necessary [3].

The Rx interface allows the sNMG to pass on required QoS characteristics in form of SDP information of the service or any other available application information to the PCRF in the SAE. This information can be, for example, the required bandwidth, media type, required priority level, etc. The PCRF then decides on appropriate internal QoS service level to be applied for this service and replies back to the sNMG regarding the QoS level the SAE can provide.

Additionally, the sNMG can request to the PCRF via this Rx interface for status report of the connection, the type of access network the terminal is using (GPRS, UMTS, etc) and when there is any change on the QoS level. This information can be exploited by the ADM in the sNMG for adaptation of the stream to each particular user.

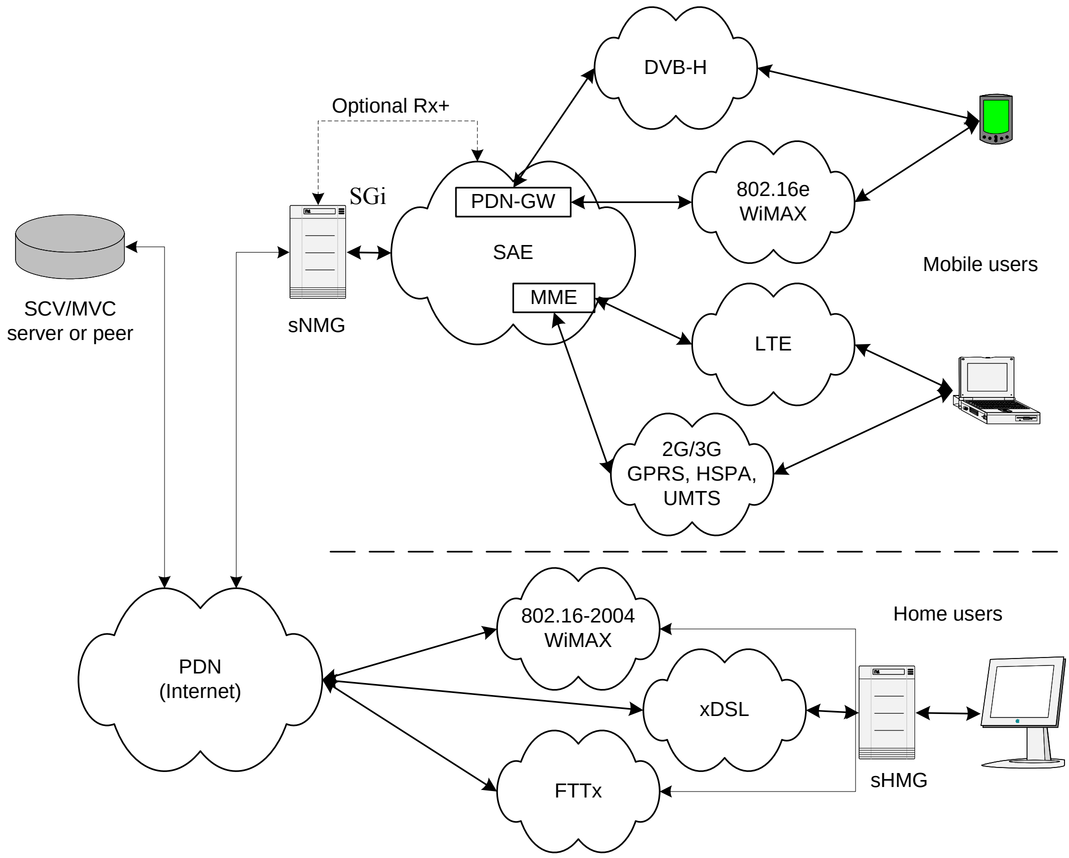

Figure 2: Integrated SEA and SAE network architecture

As one of the important network entities introduced in SEA, it is important to estimate the requirements and expected capabilities of the sNMG modules. This is essential with regard to infrastructure planning and capital investment estimation for operators that intend to offer SEA services to their customers. This document gives detailed estimation on the expected hardware resource requirements of the sNMG including some example figures of potential customer bases and resources needed to provide SEA services to them.

2 sNMG system requirements

Similarly to the sHMG , sNMG will be located at the edge of digital data transmission networks (e.g. 3GPP SAE, WiMAX, xDSL etc.). As shown in Figure 3, sNMG will have interfaces with:

a) a number of service providers,

b) multi-RAT network and

c) the IP network that interconnects the sNMG with the sHMG and the home terminals.

In the longer term, sNMG may be integrated with the IMS element, or the sNMG functionality may be overloaded on an IMS. However, building on top of an IMS would further complicate implementation and SEA might loose its focus. Thus, for simplicity reasons, within SEA, we keep IMS and sNMG as independent network elements.

Apart from interfacing the SEA architecture, the sNMG will provide the network content reconstruction and forwarding, thus enriching the delivered content quality. Within sNMG, only service provider data will be stored as local cache. However, sNMG will have the capability to identify the terminal processing, decoding, networking and display capabilities and, based on user preferences and contract, may initiate streaming video adaptation or reconstruction (thus, resulting in the enrichment of the network's QoS) before forwarding to the end user.

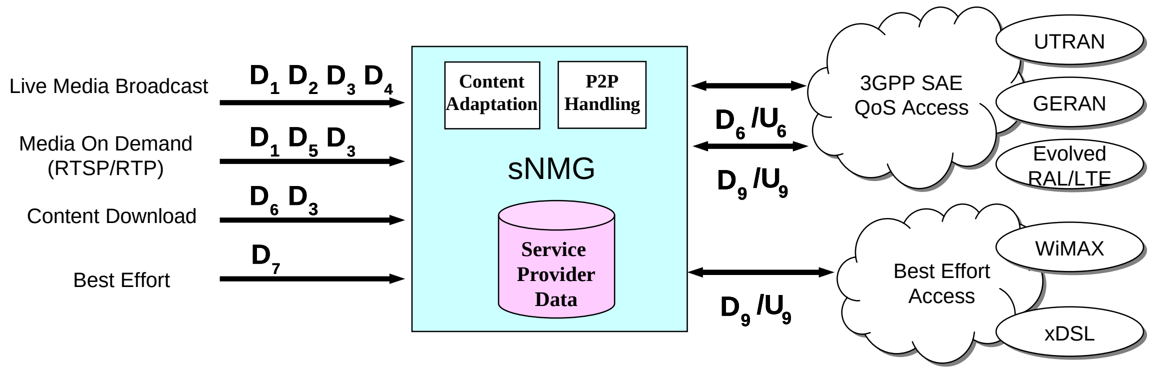

Figure 3: sNMG network architecture

2.1 sNMG interconnection bandwidth requirements

This section describes external interfaces requirements between a sNMG and other network entities.

- ***Towards sHMG.*** Since the sHMG is meant to support a small number of users, e.g. only family members within a single household, the amount of bandwidth needed for this interface is assumed to be minimal. Typically, this should be in an order of less than 10 Mbps for the downlink direction per household. For the uplink direction, e.g. from the sHMG to the sNMG, the required bandwidth is expected to be less. A typical DSL connection per household would conveniently meet this specification.
- ***Towards other sNMG.*** The communications between two sNMGs is inherently large due to the fact that both need to support large amount of users. Based on our estimation in the following section 3.1 and 3.2, a single sNMG would be able to handle some hundreds of streams concurrently. Taking into account bandwidth requirements from peer-to-peer SEACast communications and other SEA-related control information, this interface should be able to support up to several hundreds of Mbps in both directions.
- ***Towards SAE.*** The sNMG nodes placed at the border of the SAE core network as shown in Figure 2 must handle SEA traffic, both in the data sessions and control sessions (in case the

optional Rx interface to the IMS PCRF module is deployed) for all mobile subscribers. Similarly to the interface between sNMG's, this interface between a single sNMG and the SAE must handle aggregated bandwidth consumption from hundreds of users in the user plane. However, the optional Rx interface for the control plane is not expected to consume much bandwidth in comparison to the user plane interface since it is used only infrequently for QoS management and connection status reporting.

Nevertheless, the total bandwidth requirement on this interface would be the same as for the interface toward other sNMG's, up to several hundreds of Mbps in the downlink path, i.e. from the sNMG to the SAE. The bandwidth requirement for the uplink path, i.e. from the SAE (the mobile terminals) to the sNMG is expected to be lower since the mobile terminals will more likely be content consumers, and not distributors nor be participating in peer-to-peer sessions. This interface could be realized with, for example, a Gigabit Ethernet interface.

2.2 Software modules requirements

Similarly to sHMG, the SEA sNMG system hosts the following modules and communication interfaces. More details may be found in [4]. As shown in Figure 4, the software modules are identical to the sHMG except that the sNMG is not connected to any TAM and the scaling of the modules that have to handle lager amount of traffic than the sHMG. In more details the following communication interfaces and languages have been defined:

- **ADM<->NAM.** Messages in XML format are exchanged. The messages are based on MPEG21-DIA metadata specification.
- **ADM<->ADM.** Messages in XML format are exchanged.
- **RTSP Client<->Server.** SDP Messages are exchanged. SEA has defined extensions to the standard SDP messages to support different types of streams e.g. SVC, MDC, MVC.

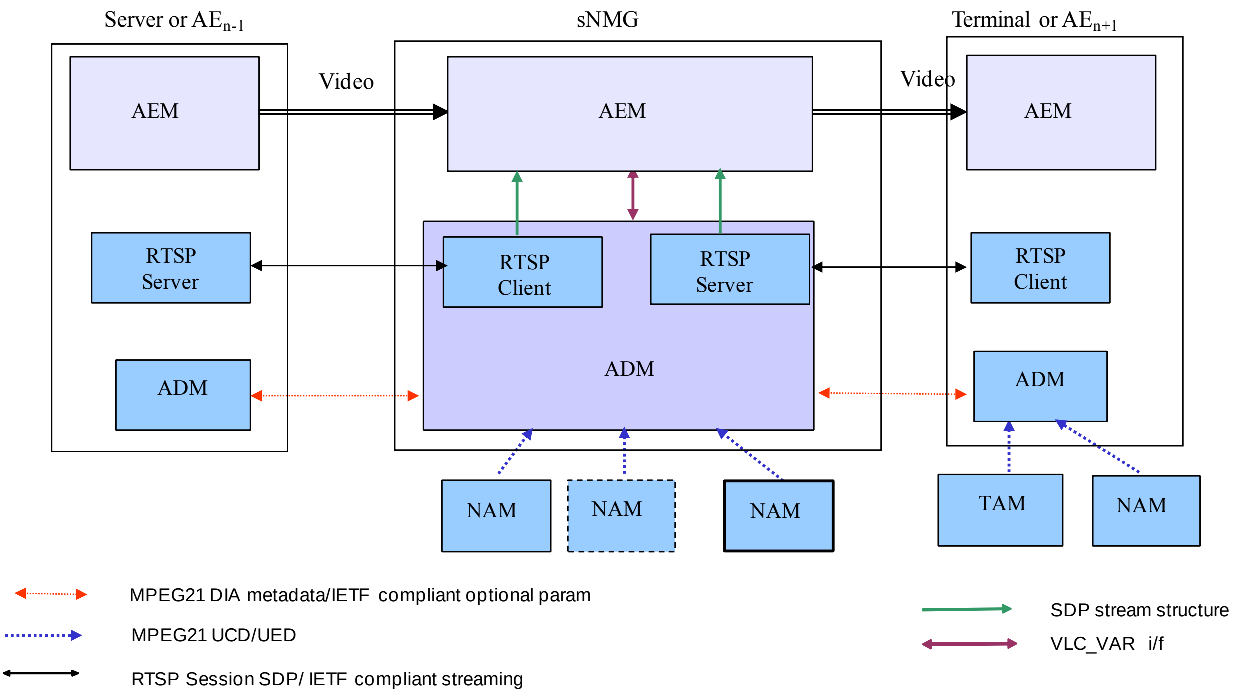

Figure 4. SEA sNMG modules communications

- **ADM (RTSP Client/Server)<->AEM.** From the SDP messages, SDP stream information is extracted. This information in the form of SDP streams structures, as already defined in VLC are exchanged between the ADM and the AEM, as shown in the figure.

- **ADM<->AEM.** Information between the ADM and the AEM is exchanged in the form of VLC shared memory variables (VLC_VAR_variable-name variables).
- **SEACast P2P<->ADM**. The SEACast P2P modules are based on the VidTorrent protocol, properly modified to meet the SEA requirements. When the SEACast terminates the P2P session and communicates with the ADM, then it operates as an RTSP Server. They exchange XML messages based on MPEG21-DIA standard with some proprietary tags.

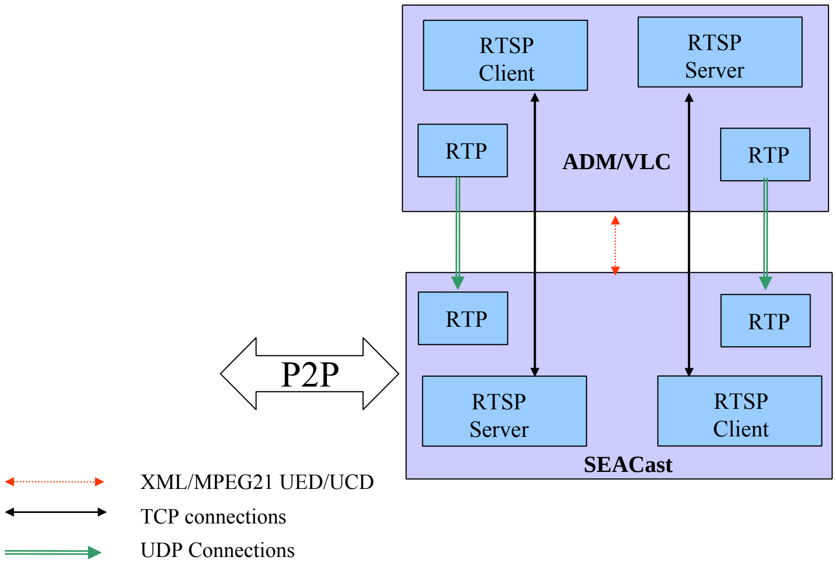

Figure 5. ADM/VLC SEACast Interface

This section describes the requirements of the different modules that are implemented in the sNMG. Some of the modules (ADM, AEM) are created as VideoLan Client plug-ins and live in the client that adapts the media data and forward the sessions to the final terminal.

Other modules are daemons that reside in the system background. The TAM, NAM and P2P applications continually monitors the status of the terminal and the networks to which the sNMG is connected, sending updates to the ADM from time to time.

In the following sections modules will be analyzed to understand their requirements. In general the involved computational complexity is reduced for adapting standards as SVC and MVC thanks to the organization of the information that allows an efficient browsing and discarding of un-required data.

2.2.1 ADM

The *Adaptation Decision Module* (ADM) is a module contained inside the VLC. This module implements the decision-taking engine that decides the data forwarding and adaptation to the receiver terminals. The ADM needs to collect and store information coming from the different modules to which it is connected. Most of the messages are in XML format. Thus, XML parsing and merging is needed.

Moreover, the ADM is responsible for the streaming connection set-up and handling. This is achieved via the RTSP protocol. For each session, an RTSP client-server pair is created and some additional information is stored (e.g. video type, adaptation rules etc.). This processing is more demanding than the XML parsing.

The computational complexity of the adaptation decision and the connections handling are the major ADM limitation factors. Additional requirements reside on the memory requirements to store the data for each terminal, network and session that it manages. The estimation of session requirements is based on the data that the module requires to store.

2.2.2 AEM

The *Adaptation Execution Module* (AEM) is a module contained in the VLC. This module realizes the adaptation of the incoming streams to the terminal and network bandwidth characteristics. Within SEA we have selected SVC and MVC video-streams. The adaptation is based on the characteristics of the video-streams (dropping/combining layers or views), thus no transcoding takes place. In this way, the procedures performed by the AEM do not require high complexity computational resources per stream.

Regarding the memory resources, AEM needs some buffering due to the fact that it may terminate peer-to-peer sessions, collect multi-layers session in a single stream to be forwarded to the terminal or stream the layered session to the final terminal.

2.2.3 NAM

The *Network Awareness Module* (NAM) has knowledge of the physical characteristics of the network (multiple access, QoS classes, coverage). Moreover, it may be able to measure or probe network parameters e.g. number of users, available bandwidth, etc.

At the sNMG level, there will be one NAM per available network interface. At the prototype this will be the NAM of a Gigabit Ethernet. Although the NAM is still in integration phase in which its hardware requirements and processing load cannot be measured at the time being, it is expected, however, that the NAM will not require significant amount of resources both in terms of processing power and memory space. This is due to the fact that the NAM only has to regularly monitor some few parameters for network conditions and reports back to the ADM once significant changes occur.

2.3 Video adaptation requirements

Transcoding a stream to a different bit rate or a different spatial resolution is in general a very complex task and requires a careful analysis of the hardware and software requirements to deal with the associated complexity of the solution. In general, reducing the bit rate or the spatial resolution of the image requires the complete decoding of a sequence and the subsequent re-encoding of all the decoded data. Such a solution becomes quickly unpractical if multiple different streams have to be transcoded in parallel. In this case, a complex mixed hardware and software solution can deal with the computational complexity problem, fixing some constraints on the maximum number of transcoded streams that can be handled in parallel. Especially, in case of sNMG this is a very difficult problem.

In SEA however, as also explained in [4], we concentrate on SVC, MVC and MDC formats. These formats are scalable, thus adaptation is achieved without any processing hungry transcoding process. As SVC and MVC are thought as annexes that extends the functionalities of the H.264/AVC standard, they inherit the data organization from their predecessor. In particular, atomic information is contained in a data unit called *Network Abstraction Layer Unit* (NALU). The NALU is the basic constituent of a H.264 bit stream. The NALU can be easily parsed as the first bytes contain some basic information about the contents of information of the NALU. A parser can easily identify the contents of the NALU and decide if it needs to keep it or discard it depending on its particular purposes, thus reducing greatly the computational power requirements for this process. This greatly improves the performances when considering some bit stream transcoding or adaptation.

The following sections deal with computational complexity, memory and storage requirements for the different codecs considered for the SEA project, in order to assess the minimal requirements for the prototype implementation.

2.3.1 MultiView Video Coding requirements

MultiView Video Coding (MVC) bit streams data are organized in NALUs that separates the different views in different packets, so that they can be easily parsed and un-required views may be easily discarded. Views are identified through a field in the header that allows selecting the required views, so adaptation is so easy as parsing that number and decide if we should keep the packet or not.

Some complications resides in the fact that some views may rely on other views for decoding, so the number of forwarded views may be higher than the effective number of views selected by the system for decoding and displaying.

2.3.2 Scalable Video Coding requirements

In Scalable Video Coding a NALU contains information pertaining a specific layer of spatial, temporal and quality scalability. The basic layer (called base layer) contains information encoded in the original AVC format. The NALU header contains details about the scalability level the NALU under examination is pertaining to. The D, T and Q triplet is a set of integer values that identifies respectively the spatial, temporal and quality layer of the NALU. For example a triplet {0,0,0} identifies an AVC-encoded base layer NALU, while some integer positive numbers for any of the three components identifies the NALU for one of the possible enhancement layers.

Shaping the bit stream according to the characteristics of the terminal or the available bandwidth is therefore very easy. The sNMG is informed by the client or the sHMG about the receiver's display capabilities and is aware of the available network capacity; hence it can decide for a target capability layer that fulfills all the constraints. Given this decision the sNMG should parse the incoming NALUs and decide whether to forward or drop the NALU according to the D, T and Q triplet and the decision it took.

2.3.3 MDC requirements

Multiple Description Coding (MDC) allows a more robust transmission of information sending different representations of the same sequence through different channels. If one representation is received at the receiver side, a basic quality decoding is performed, but if all the representations are received, the extra data allows improved performance in terms of quality. The trade-off in MDC is between redundancy (correlated information is delivered through different channels, hence the encoder does not completely remove the correlation in the video signal) and the quality in case of errors (the capacity of the MDC scheme to recover the missing information).

The solution we provide in the project is based on *Redundant Slices* (RS). RS are introduced in the H.264/AVC standard as a means of providing some extra information in case some of the original encoded information is not received correctly. We used this technique as a means of differentiating the information sent over different channels, as primary and redundant slice representations are alternated in the descriptions. One great benefit of this approach lies in the fact that redundancy may be finely tuned through a rate-control that decides how many bits can be allocated for the primary and the redundancy representations.

The optimal amount of redundancy (i.e., the QP value of the redundant representation of each slice) can be worked out analytically as a function of the total available rate, the probability of packet loss, the video RD characteristics and the position of the frame within the GOP (which affects the amount of drift introduced in case a primary representation is replaced by its redundant counterpart). Hence, the total redundancy introduced by the MDC scheme depends on the total rate, the video characteristics and the packet loss rate. As a rule of thumb, optimal redundancy levels for $p = 10\%$ are about 30-35% in a large range of total rate values, about 20-25% when $p=5\%$ and 15-20% when p is as low as 1%.

MDC decoding requires minimal complexity at the receiver side. RS are marked in the NALU header and can be easily identified parsing the first bytes in the NALU header. When RS are decoded and the corresponding primary information is received as well, we can drop the redundant information and decode the image at full quality detecting the slice type with minimal computational effort.

MDC encoding requires generating the redundant slices. This part requires encoding the scene with a H.264/AVC encoder and generating extra data for the RS layer. As the incoming bit stream does not contain RS information in general a transcoding (a decoding with an appropriate encoding and generating RS) is required. This requires a significant computational complexity.

As a conclusion, within SEA, the video adaptation requirements are minimized. Yet, scaling issues (large volume of terminal users) will increase the requirements and may require a parallel or cluster based sNMG architecture.

2.4 SEACAST requirements

The SEACast P2P engine is hosted on the sNMG, but it resides as an independent module and not as another VLC plug-in. In the SEA described scenarios, the P2P sessions may be terminated at the sNMG and data is forwarded to the final terminal. VLC in all these situations acts as an adaptation engine. Both the VLC and the SEACast resides on the same machine. The transmission between the two elements occurs via a localhost IP port.

Information about media characteristics are exchanged between the ADM and the SEAcast. For the SEAcast implementation, we have extended the VidTorrent engine (implementing innovative algorithms for tree creation, implementing multiple description propagation along the trees, solving bugs in the code in order to make it stable, scalable to a few hundreds of clients, etc). From preliminary simulations over the PlanetLab test bed, SEACast has shown the capability of scaling up to the maximum possible extent (i.e., over all the available nodes).

The SEAcast P2P application is based on a modified version of the VidTorrent protocol. It generates an overlay network that exhibits a multiple-tree architecture. This topology is particularly suitable for MDC distribution, as descriptions can be independently propagated along different trees in the overlay. In its present implementation, SEAcast manages two trees and two descriptions, generated as explained in Sec. 0. Hence, a traffic overhead due to the use of MDC has to be taken into account.

The MDC scheme depends on the total rate, the video characteristics and the packet loss rate. As a rule of thumb, optimal redundancy levels for p = 10% are about 30-35% in a large range of total rate values, about 20-25% when p=5% and 15-20% when p is as low as 1% [8].

2.5 Storage requirements

As already described, another innovation in SEA, also related with the foreseen business model and the SEA approach to provide on-the-fly enriched PQoS, is that sNMG will be able to offer multimedia content adaptation and caching/storage functions to the service provider and the community.

The storage capacity of the sNMG will be used only for service provider and/or the community network caching. A/V content files, layers or segments will be stored and indexed there, adapted and retrieved/relayed on remote (guest) users or subscribers' request. In this way, P2P (super) distribution and load balancing will also be feasible.

3 Preliminary testing for the hardware requirements of the sNMG

In this section, we give estimations of the hardware resource requirements of the sNMG to support a given amount of users.

The different components of the sNMG as described in section 2 would all contribute differently to the total hardware resource requirements. From a first analysis the modules that will significantly contribute to the sNMG load will be the SEACast module and the MDC encoding. Its resource requirements such as CPU usage and RAM are estimated to be proportional to the stream's bit rate, e.g. the higher the bit rate, the more media segments and connections they have to handle. On the other hand, the ADM and the AEM are not expected to increase proportionally the CPU and memory load, as the ADM refers mainly to a network-terminal pair and not only to terminals, while AEM mainly handles scalable video content.

This section gives a very draft approximation for an upper limit of supported users with a given set of hardware specifications. The test setup as depicted in Figure 6 shows a server streaming videos to N clients via the sNMG. There are N VLC applications running at the sNMG, which are configured simply to relay the streams to their destinations. This is comparable to the case where the SEACast is not used and the adaptation done by the AEM implies simply dropping entire sessions or no dropping at all, but not dropping individual packets from a session.

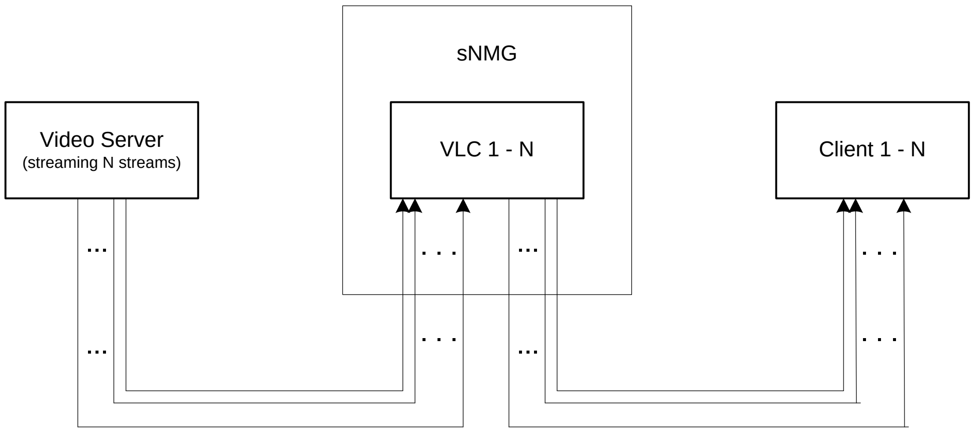

Figure 6: test setup for the sNMG load test

The hardware used for the sNMG has 1 GB of RAM and a dual-core CPU of 2 GHz each. There are two MPEG-4 test streams to be used with bitrates of 2 Mbps and 2.7 Mbps. The experiments were done in two main sets, one with the 2 Mbps stream and another one with the 2.7 Mbps stream. In each set, the number of streams passing through the sNMG was set to be either 1, 4, 8 or 12 streams, meaning there were 1, 4, 8 or 12 VLC applications running at the sNMG, respectively.

Table 1 shows the CPU usage in percent measured at the sNMG for different number of streams and at two different bit rates. Note that the column labeled "SEA average CPU usage" represents the average CPU usage without other background CPU usage which was measured to be approximately 5.01%.

Number of streams	2.0 Mbps stream				2.7 Mbps stream			
	CPU 1 usage (%)	CPU 2 usage (%)	Average CPU usage (%)	SEA average CPU usage (%)	CPU 1 usage (%)	CPU 2 usage (%)	Average CPU usage (%)	SEA average CPU usage (%)
1	8.93	1.17	5.05	0.05	9.28	1.40	5.34	0.34
4	11.28	2.45	6.86	1.86	12.75	3.48	8.12	3.11
8	12.00	5.18	8.59	3.59	12.26	6.60	9.43	4.43
12	12.65	6.82	9.74	4.73	12.35	12.05	12.20	7.19

Table 1: CPU usage for different numbers of streams

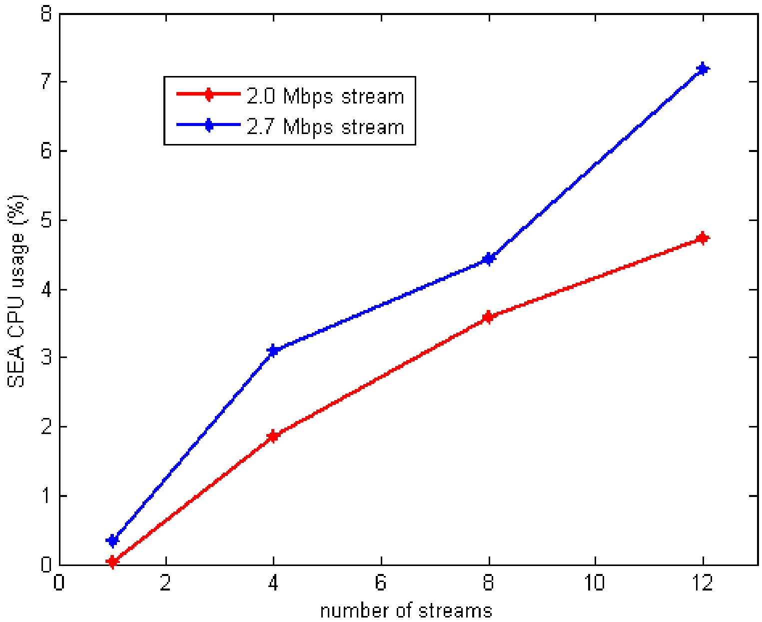

Figure 7: SEA CPU usage for different numbers of streams

In Figure 7, the SEA CPU usage is shown in relation to the bitrates and the number of the streams. Although each additional stream increases the CPU load by less than 1%, this implies that there is a certain limit in the number of streams the sNMG can support simultaneously.

As an example in our case, if the relationship between the CPU consumption and the number of streams is assumed to be linear, an additional stream would require 0.34% and 0.60% of CPU power for 2.0 Mbps and 2.7 Mbps streams consecutively. Given that there is c.a. 95% of the CPU power left excluding all the background processes, this sNMG could simultaneously support no more than either 280 streams of 2.0 Mbps bitrate or 167 streams of 2.7 Mbps bitrate.

Since the notion of CPU usage in percentage is dependent on the CPU power of the hardware, Table 2 converts and shows the CPU requirement to support a stream with different bitrates in terms of processing cycles of the CPU instead. Hence, one can roughly estimate the upper limit on the number of supported streams of any hardware given that its CPU power is known. Additionally, the value for a 1.0 Mbps stream is approximated from the data for the 2.0 and 2.7 Mbps streams.

	1.0 Mbps stream	2.0 Mbps stream	2.7 Mbps stream
CPU usage (MHz)	3.80	6.80	12.00

Table 2: CPU resource requirement to support a stream with different bitrates

Note that this estimation is done using a typical PC as the sNMG. In reality, the performance could be better with specialized hardware, e.g. the total number of supported streams concurrently could be larger.

3.1 Transrating and Transcoding effort estimation

SEA clearly does not address the case of transrating or transcoding. The only video streams that are handled (streamed and adapted) are the H.264/SVC and the H.264/MVC types. However in a commercial deployment situation, transrating and transcoding may take place in the network (either at the sNMG or at other similar servers). Thus for completion reasons, in this deliverable, we append some results of transrating and transcoding that were generated using the standard VLC library (as it is included in the normal VLC distribution)

The test setup is the same as the one depicted in Figure 6. It shows N servers streaming videos to N clients via the sNMG. But sNMG is configured to relay the streams to their destinations changing the bitrate of the video. The hardware used for the sNMG has 1 GB of RAM and a dual-core CPU of 2 GHz each.

3.1.1 Transrating

Figure 8 describes a 2Mbps MPEG2-TS H264 source streams. The experiments were done in 3 main sets, one with a transrating ratio of 1/2, 1/4 and 1/8. In each set, the number of streams passing through the sNMG was changed from 1 to 2 streams due to CPU charge limitation.

Note that in order to get a 2 Mb/s stream in the network (send rate), we had to set the VLC program to a 4 Mb/s data bitrate.

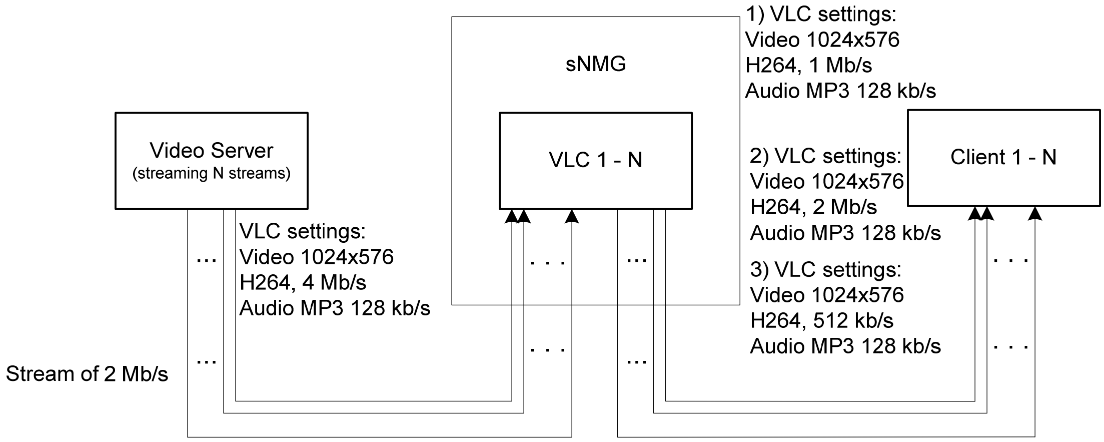

Figure 8: test setup for the sNMG Transrating

Number of streams	Transrating of 1/4	Transrating of 1/2	Transrating of 1/8
	SEA Average CPU usage (%)	SEA Average CPU usage (%)	SEA Average CPU usage (%)
1	45,45	45	36
2	81,50	88	71,6

Table 3: CPU usage for different type of transrating

Table 3 shows the CPU usage in percent measured at the sNMG for different type of transrating. Note that the tests have been done with a 1024x576 pixels video with 24fps with a 5334Mb/s original bitrate. Video Streamers are VLC servers encoding the original video towards a H264 video encoded in 4 Mb/s and 2 channels MP3 audio encoded in 128 kb/s in order to get an observed stream of 2Mb/s. Other tests to transrate a 1920x1200 video has been done but the streamer server has not enough CPU resource to stream it in a H264 stream.

Notes that the router CPU charge does not apparently affect the routing performance. The router with a 90% CPU usage still routes 64 Mb/s of data without creating errors or dropping packets.

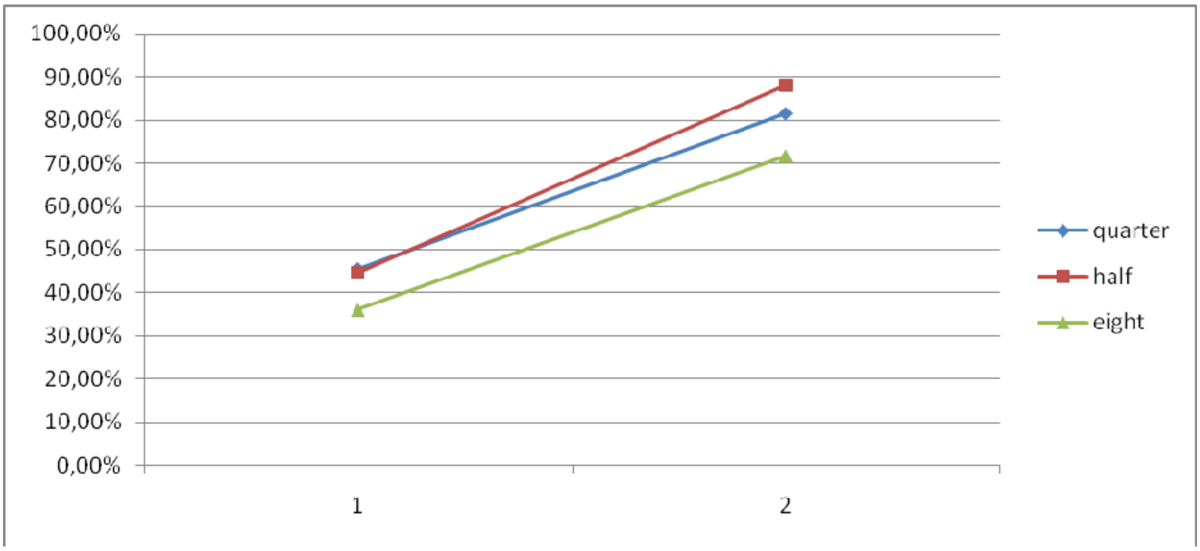

Figure 9: SEA CPU usage for transrating of 2Mb source stream

3.1.2 Transcoding

The experiments were done in 2 main sets, one transcoding towards a 2Mbps DIV3 stream and another one to a 2Mbps WMV2 stream. In each set, the number of streams passing through the sNMG was changed from 1, 4 and 8 streams, meaning there are also 1, 4 and 8 VLC applications set to transcoding scenarios.

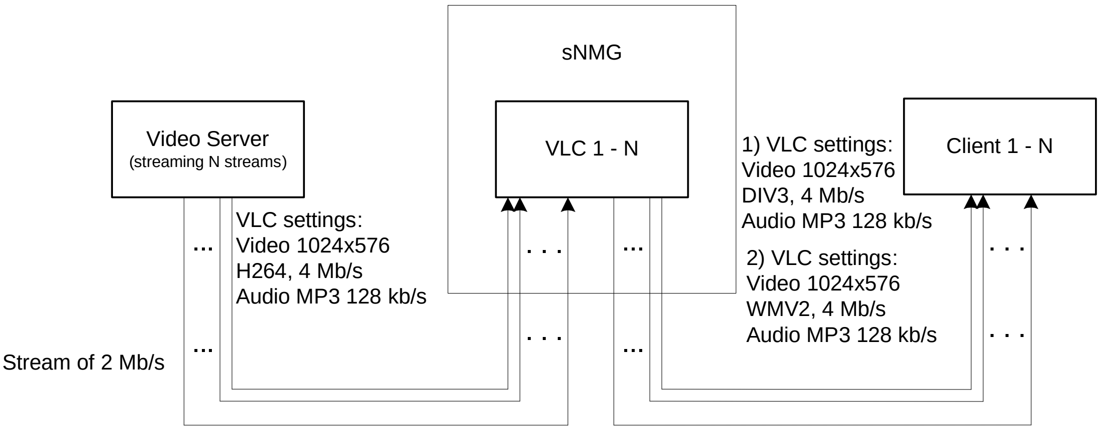

Figure 10: test setup for sNMG transcoding

Number of streams	Transcoding to DIV3	Transcoding to WMV2
	SEA Average CPU usage (%)	SEA Average CPU usage (%)
1	12,20	
2	26,50	25,83
3	40,11	46,56
4	52,93	
5	66,05	
6	78,35	78,35
7	90,15	
8	99	98

Table 4: CPU usage for different types of transcoding

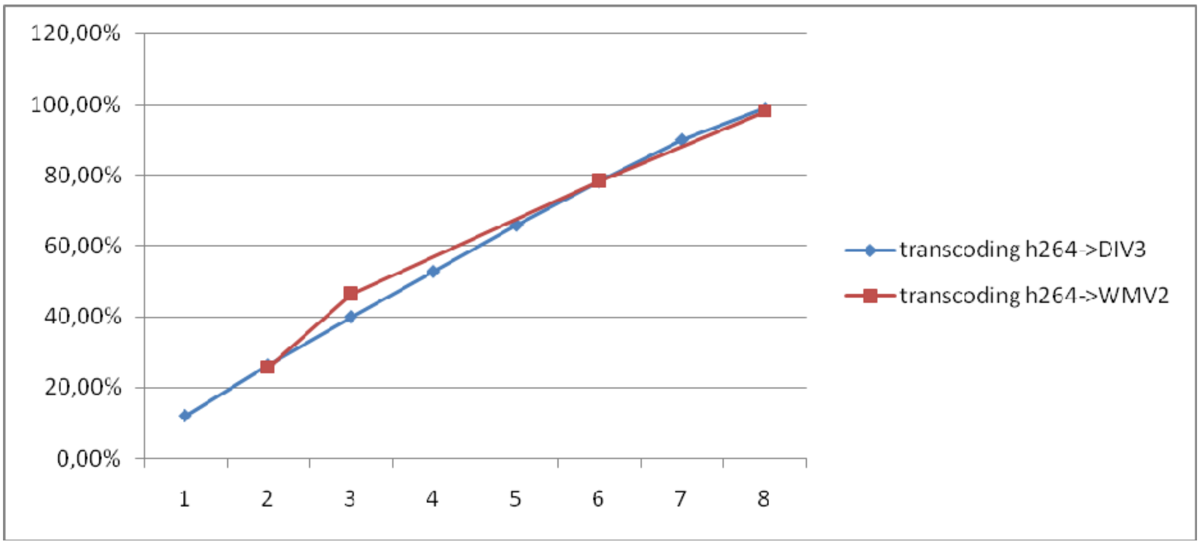

Figure 11: SEA CPU usage for Transcoding of H264 transcoding to DIV3 and WMV2

Other test regarding media traffic impact in a congestion scenario can be seen in the following Figure. The tests involve high FTP and 2 Mb/s RTP traffic and results in erroneous packets in the FTP traffic but limited packet loss in the RTP traffic. This test takes into account the transrating and transcoding, but just the scenario in Figure 10.

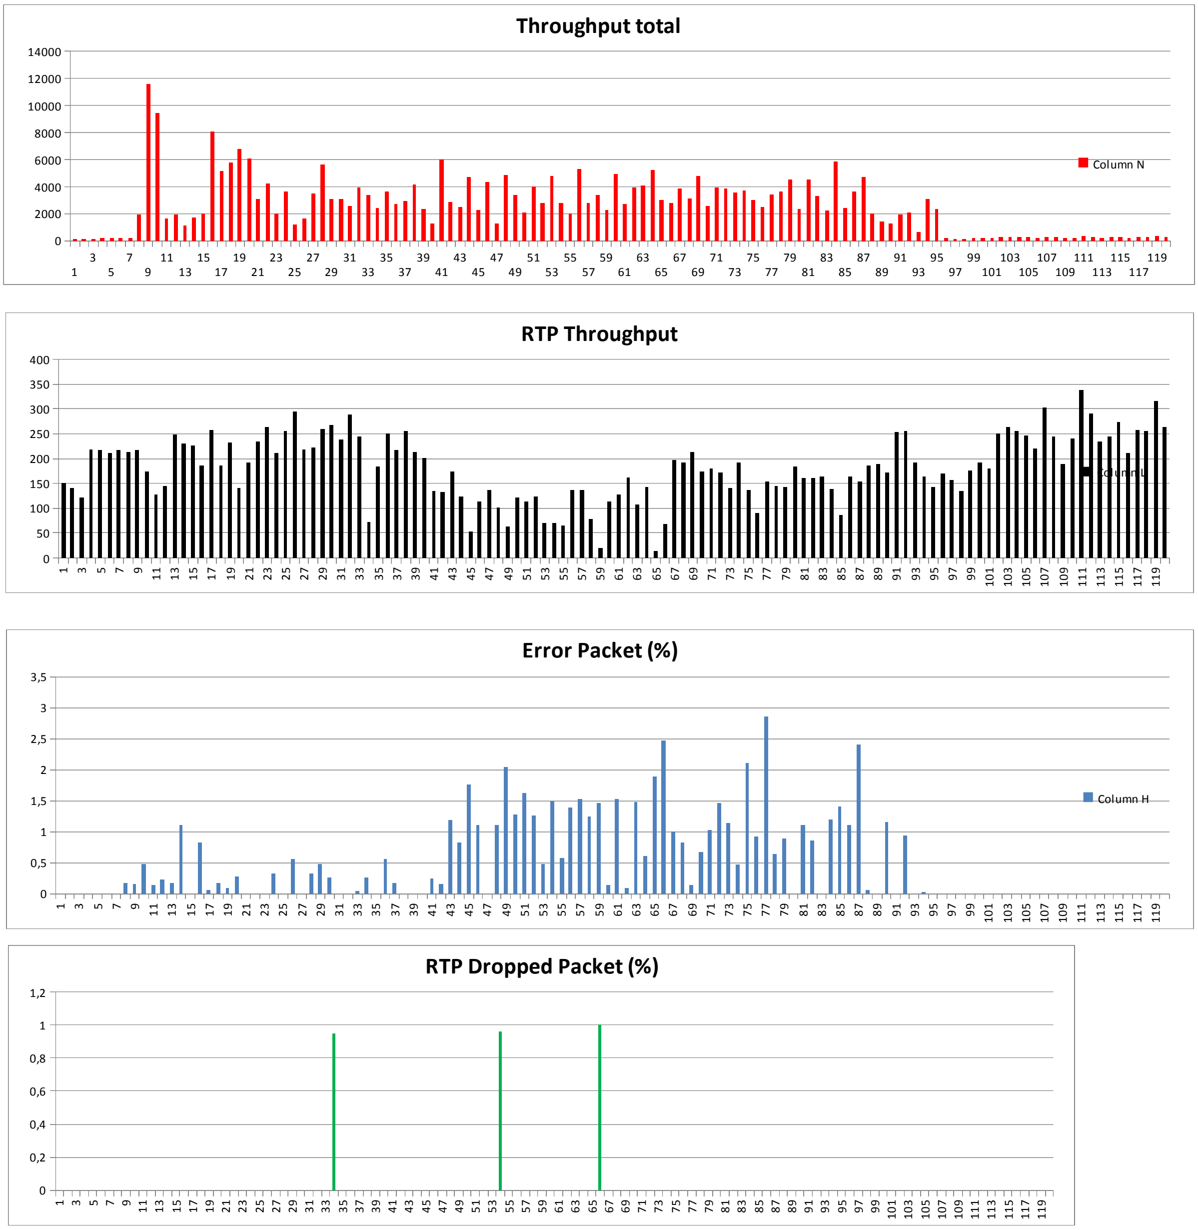

Figure 12: Traffic impact in a congestion scenario

3.2 SEACast module

We performed different runs of simulations basing on different conditions. The first condition to be tested is sharing a self-produced multimedia content to the network. Secondly, we wanted to try the effort required to run a rewire (a server that enables connections between different peers, acting as a router) and finally we tried joining a tree and receiving a multimedia content. We tried sharing different multimedia contents with different resolutions and different bitrates to test the memory requirements in different situations.

The data were obtained through the OS data structures. In particular, Linux OS has a special folder named /proc where it saves all statistics about running processes. The information is stored as text format and can be easily browsed.

```
State:  S (sleeping)
Tgid:   4591
Pid:    4591
PPid:   4590
TracerPid:      0
Uid:    1000    1000    1000    1000
Gid:    100     100     100     100
FDSize: 256
Groups: 16 33 100
VmPeak:    55284 kB
VmSize:    55284 kB
VmLck:         0 kB
VmHWM:     29512 kB
VmRSS:     29512 kB
VmData:    18760 kB
VmStk:       132 kB
VmExe:         4 kB
VmLib:     30960 kB
VmPTE:       104 kB
Threads:        2
SigQ:   3/8188
SigPnd: 0000000000000000
ShdPnd: 0000000000000000
SigBlk: 0000000000000000
SigIgn: 0000000001001000
SigCgt: 0000000180004ea
CapInh: 0000000000000000
CapPrm: 0000000000000000
CapEff: 0000000000000000
voluntary_ctxt_switches:        734
nonvoluntary_ctxt_switches:     393
```

Figure 13: Receiver process statistics (*/proc/4591/stats* file)

Figure 14 shows the SEACast CPU usage in relation to both the video bit-rate and three different video formats (QCIF=176x144, CIF=352x288, PAL=720x576) on a Centrino Duo 1.8 GHz machine equipped with 2 GBytes of RAM.

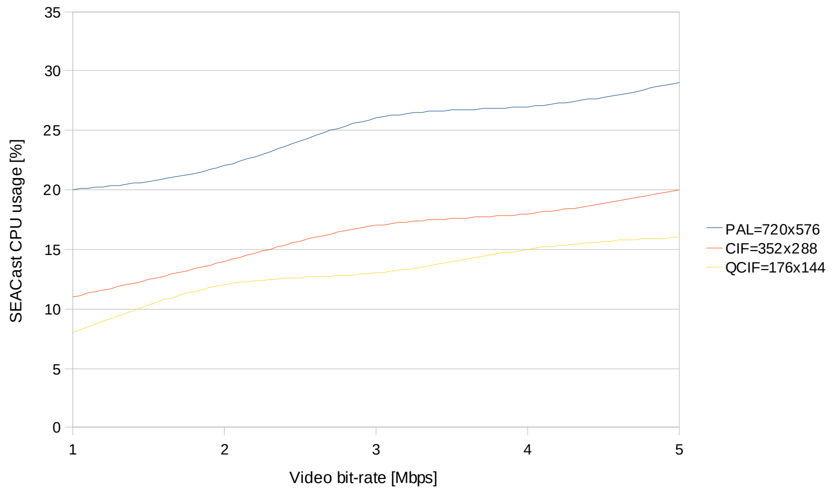

Figure 14: SEACast CPU usage

Figure 13 shows an example based on the collected data. Among all the available information, virtual memory considers the total process allocated memory, consisting of the program and the data application memory. In particular the Peak Virtual Memory data allows inferring required data for each process.

4 Scaling & Network Dimensioning

In this section we provide a draft subscribers and network dimensioning estimation. It has to be mentioned however, that emphasis has been put on the system/prototype functionality and not on scaling issues, which are mainly related to deployment parameters.

4.1 From broadband access network operators' perspective

This section discusses a possible deployment scheme of the SEA sNMG by broadband network operators. These operators can be, for example, mobile network operators or land-line DSL providers who wish to provide SEA services to their customers.

In this case, the operator owns and deploys the sNMG as a part of their network. The amount of hardware resources, i.e. how many sNMG modules to install, bandwidth to be allocated, etc all depend on various factors such as marketing strategy, the size and type of the customer base of each operator. Most of the DSL operators' customers might be, for example, household or SME users where SEA service and the sHMG might be offered bundled with the DSL modem in reduced price. This can possibly increase the penetration rate of the SEA service while mobile users of a mobile operator can only enjoy partial services of SEA which might result in lower subscription rate.

In the following, we give an example for infrastructure dimensioning to support SEA services by a mobile operator in some European markets, we will estimate the number of SEA subscribers in these markets first. Reasonable penetration rate for SEA could be approximated roughly based on the first acceptance phases of mobile TV today.

From the results of Finnish Mobile TV [5] project[1], 58% of the participants believed that mobile TV will become popular in the future while 41% of these expressed an interest in buying the service. The reader should also take notice here that the Finnish market is considered one of the most matured European markets, where large proportion of the population is familiar with the notion of video streaming on mobile terminals.

Another online market research by the Forschungsgruppe Medien GmbH, which was conducted recently and which reported on the current intention of German mobile subscribers in continuing or increasing their use of mobile TV services, indicated that 20% have answered positively.

Although it would be unrealistic to assume one common penetration rate to be applied to all European countries, as that is much subject to the cultural diversities and the distinctions between the network technology investment strategies followed by the operators in each country, however, we could still assume a penetration rate of 20% of all 3.5G subscribers (i.e. having a downlink speed of ≥3.6Mbps), as a rough estimation to get the picture of how much invested resources would be needed t for the commercial launching of SEA services.

As indicative examples about the targeted subscription bases in a European medium-sized and larger-sized mobile telecom market, we take into account the number of 3.5G (i.e. HSDPA) connections in the Vodafone subsidiaries of Greece and Germany, respectively (Table 5 and Table 6). Similarly, so as to escalate the estimations to a single service provider in a pan-European market of subscribers, we consider the equivalent figures for the European Vodafone Group, as a whole (Table 7).

[1] The Finish Mobile TV was an extensive two-year project for pilot operation of mobile TV services that gave detailed reports on Finish people's acceptance about the delivery of video programs (either online or on-demand) on their mobile handsets.

Year	2008	2009
total connections	5,465,465	5,804,546
Number of WCDMA family connections	1,118,382	1,457,216
% of WCDMA family connections	20.46%	25.10%
% of WCDMA-2000 (384 Kbps) connections	16.30%	16.45%
% of WCDMA HSDPA (7.2/14.4 Mbps) connections	4.17%	8.65%

Table 5: Subscriber statistics of Vodafone Greece

Year	2008	2009
total connections	34,412,000	36,361,046
Number of WCDMA family connections	5,836,000	8,505,219
% of WCDMA family connections	16.96%	23.39%
% of WCDMA-2000 (384 Kbps) connections	13.14%	15.88%
% of WCDMA HSDPA (7.2/14.4 Mbps) connections	3.82%	7.51%

Table 6: Subscriber statistics of Vodafone Germany

Year	2008	2009
Total connections	180,066,145	187,545,099
Number of WCDMA family connections	32,061,015	46,726,184
% of WCDMA family connections	17.81%	24.91%
% of WCDMA-2000 (384 Kbps) connections	14.06%	17.46%
% of WCDMA HSDPA (7.2/14.4 Mbps) connections	3.75%	7.45%

Table 7: Subscriber statistics of European Vodafone Group

If we apply the 20% penetration rate for SEA services to the number of subscribers with 3.5G-capable terminals in Germany in 2009, for example, there would be approximately 550,000 subscribers to SEA. If we further assume that 10% of all these SEA subscribers would use the streaming services at any given time simultaneously, Vodafone Germany would have to invest enough hardware resources to support 55,000 clients. If the average bit rate for a stream is 1 Mbps, this translates into a total required CPU power of 209 GHz using the data from Table 2.

In reality, the required CPU power will be much higher, as this responds to homogeneous distribution of users through out the network.

4.2 From CDN providers' perspective

Another deployment scenario for the SEA sNMG would be by the Content Delivery Network (CDN) providers. These companies, such as Akamai [5] or Limelight Networks [6], provide infrastructure services to efficiently host and distribute contents over the internet by means of deploying hundreds or thousands of edge servers at different geographical and logical areas. Thus, the contents can be cached and distributed to servers closer to the users, reducing latency and congestion in the network and offering services to larger customer base.

Since the sNMG is also a kind of an "edge server" that provides streaming services of the SEA, it can also be deployed by the CDN providers along side their edge servers to provide SEA services too. This also means that if the sNMG deployed by the CDN provider is connected or close to the SAE part of the mobile operators, it is possible for the clients to enjoy the services provided by SEA even if the mobile operators themselves do not own the sNMG in their networks.

The amount of sNMG and its required hardware resources are difficult to estimate since this depends also on, for example, the customer bases of each operator, how extensive their CDN network covers, penetration rate of SEA in the geographical areas covered by each operator, etc. Also, the information used to estimate the number of potential SEA subscribers used in section 4.1 which mainly regards mobile subscriber statistics is not applicable in this case since the subscriber base for this case could include, for example, household subscribers and other customers that mobility is not an issue. However, once the CDN provider has determined how many concurrent SEA clients it has to be prepared for, it can use information from Table 2 to calculate for required hardware resources and the amount of sNMG it has to install later on.

4.3 From independent SEA service providers' perspective using Cloud Computing concept

This is another deployment scenario for SEA sNMG where an independent service provider who does not own nor willing to invest large capital into setting up and maintaining sNMG's can still offer SEA services to customers. With the Cloud Computing technology becoming more mature these days with several Cloud Computing providers such as IBM, Amazon, etc offering there computing power in the market, deployment of SEA sNMG is not necessarily costly anymore.

One of the Cloud Computing concepts is to provide computing Infrastructure as a Service (IaaS) to the customers to run their applications and further services using distributed computing nodes. The customers then have access and control to their application running on "virtual" computing nodes in the cloud via a web interface instead. Amazon Elastic Compute Cloud [7] is one example of the Could Computing providers.

This approach eliminates the needs of the SEA service providers to invest in hardware setting up sNMG's and maintaining them, but rather paying for computing hours and used resources instead. SEA service providers who wish to start services with this approach can rent the computing resources distributed well over their target subscribers' locations, install the sNMG software on distributed nodes and start operation. Another benefit is the scalability. The SEA service providers run very few risk of under or over dimensioning their infrastructure since they can simply rent more or reduce the amount of rented hardware from the Cloud Computing providers as the number of SEA subscribers changes.

However, the sNMG deployed in this way might lack the NAM functionality due to the fact that the NAM requires probing into the lower layers of the physical connections to the sNMG nodes in which might be proved difficult by simply installing sNMG software on the computing nodes the SEA service providers don't have control.

5 Conclusions

This document has analyzed the different components of the sNMG, their expected required resources and overall requirements to realize the sNMG as a whole. Some preliminary analyses of the hardware resource requirements in terms of computing power were conducted and some example deployment scenarios of the sNMG were given. Although the accuracy of the estimation is compromised by the fact that some modules to be located in the sNMG are still being developed and not available for this study, the results still give the upper limit on the supported number of users for any given hardware.

Yet, it has to be noted that before system deployment, extensive simulation and potentially cluster or cloud-based design should be applied in order to achieve the best price vs performance vs scaling vs flexibility ratios (including also management, support, robustness, survivability and upgradeability costs).

References

[1] SEA, "D2.2: Architecture Specification," 11 July 2008, http://www.ist-sea.eu/Public/SEA_D2.2_NOM_FF_20080711.pdf

[2] SEA, "D4.1: Cross layer control, adaptation modelling and metadata specification," 30 September 2008, http://www.ist-sea.eu/Public/SEA_D4.1_SYN_FF_20080930.pdf.

[3] 3GPP, TS 23.203 "Policy and Charging Control Architecture (Release 9)"

[4] SEA Consortium, Deliverable 4.3.1 "sHMG system prototype description", May 2009, http://www.ist-sea.eu/Public/SEA_D4.3.1_PCL_FF_20080515.pdf

[5] Eija Kaasinen, Tuomo Kivinen et al., "FinPilot2 Final Report – User Acceptance of Mobile TV Services", Forum Virium Helsinki, 6 May 2008http://www.akamai.com/

[6] http://uk.limelightnetworks.com/index.php

[7] http://aws.amazon.com/ec2

[8] T. Tillo, M. Grangetto, G,. Olmo, "Redundant Slice Optimal Allocation for H.264 Multiple Description Coding, " IEEE Trans. On Circuits and Systems for Video Tech., Vol. 18, No. 1 GB of RAM and a dual-core CPU of 2 GHz each. , pp. 59-70, Jan. 2008