On-demand Regional Television over the Internet

Haakon Bryhni, Hilde Lovett, Erling Maartmann-Moe,
Dag Solvoll and Tryggve Sørensen
Norwegian Computing Center (NR)
Gaustadalléen 23, P.O.Box 114 Blindern, N-0314 Oslo, Norway
+47 22 85 25 00
E-mail: bryhni@ifi.uio.no, {Hilde.Lovett, Erling.Maartmann-Moe,
Dag.Solvoll, Tryggve.Sorensen}@nr.no

ABSTRACT

We present a regional news on demand television service over the Internet developed for the Norwegian Broadcasting Corporation (NRK), using a distributed video server and a WWW-integrated software-only MPEG-1 video player. Video is played directly over the net from the server source, thus avoiding download of a video file. As far as we know this is the first regular real-time news-on-demand service where a television station on a daily basis digitizes a full news transmission and makes it available to the public over the Internet. We describe the overall system architecture and outline the production process digitizing daily newscasts. The MPEG-1 video player developed in the project is described in some detail and the communication and audio/video synchronization issues are discussed in particular.

KEYWORDS

Network-based MPEG Player, On-line Newscast, Production Infrastructure, Video-On-Demand.

1 INTRODUCTION AND OVERVIEW

The LAVA[1] project develops a video-on-demand system based on high end workstations and broadband communication technology. The LAVA video player provides a simple VCR functionality and is integrated with a WWW browser. This client communicates with a video server on a workstation cluster [1]. The communication is implemented as a separate control and data stream, and standard video encoding formats are used.

1. LAVA (Delivery of Video over ATM) is part of the Norwegian national broadband programme HUGIN financed by the Norwegian Research Council. LAVA is a collaboration project between 5 Norwegian research institutions; NORUT (Tromsø), IDT - Institute for Data technology and Telematics (University of Trondheim), the Central University IT department (USIT Oslo) and Sintef Delab (Trondheim) with NR - Norwegian Computing Center (Oslo) as the project leader.

ACM Multimedia 96, Boston MA USA

The LAVA project also provides services based on the LAVA video-on-demand technology [2]. One such service is a local news-on-demand service developed for NRK (The Norwegian Broadcasting Corporation). As far as we know this is the first regular real-time news-on-demand service where a television station on a daily basis digitizes a full news transmission and makes it available to the public over the Internet. This paper focus on the news-on-demand service and on the LAVA video player.

In Norway, like many other western European countries, the state-owned television company is a very important channel in the media landscape. Even after years of deregulation and introduction of commercial TV and satellite channels, the state channel still has a strong position among the viewers in Norway. NRK has the most complete coverage both technically (99% of the population), contentwise and in its coverage of news.

The main strategic effort for NRK the last few years, except for the winter olympics in Lillehammer 1994, has been the introduction of regional news transmissions. Norway is a country of only 4,5 million people, but it stretches 2000 km from north to south and there are major differences between the regions concerning types of industry, culture and even language. NRK has started regional television production from seven regional offices covering the country from north to south.

This regional infrastructure is utilized so that every weekday NRK transmits local news for 15 minutes in each region. These local news transmissions consist of rather short news clips, typically one to two minutes, almost exclusively made by the NRK regional offices and covering local events.

Being based on traditional terrestrial broadcasting, these news transmissions are necessarily limited to a fixed geographical region. Local news are therefore limited by the traditional space and time limitations of television; they have to be viewed simultaneously with the transmission, and only in the designated region.

There is, however, a significant interest for these local news transmissions outside the regions were they originate. With increasing mobility in the workforce, more people live outside the region were they grew up or have their family, creating a demand for local news outside the region were it originates. In such cases it will often be an advantage that news can be watched on an on-demand basis. This makes it a very good case for experimenting with news on demand services based on Internet technology. The LAVA project therefore started cooperation with one of the northernmost of the regional offices of NRK in 1995; the office in Tromsø. The program used in this project is called "Nordnytt" (i.e. news from the north).

The infrastructure used by the local news service is the Norwegian ATM-based "Supernett" which connects the four Universities in Norway with 34 Mbit ATM connections. Video is stored on servers connected to the Supernett and the video news can be played on-line using the LAVA player described in this paper.

The video server is implemented on a cluster of workstations interconnected by ATM, using a simple Video on Demand protocol. This video server and protocol is called Elvira [3] [4] and is developed as part of the LAVA project by IDT (Institute for Data technology and Telematics, University of Trondheim). The users of the LAVA video player are given simple VCR capabilities in interaction with the news clips. Contrary to many other systems, the news video clips are not downloaded before playing, but is streamed off the server in real-time as shown in figure 1.

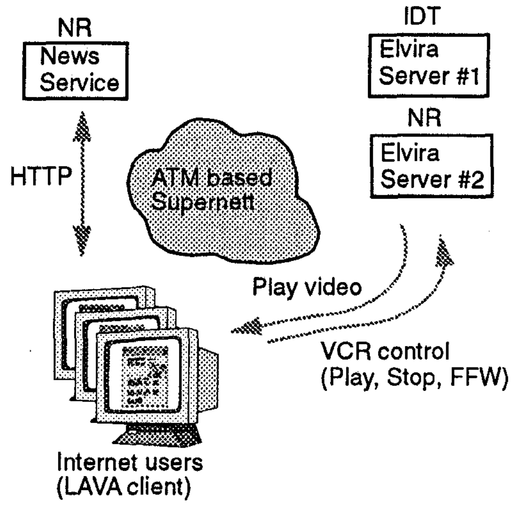

Figure 1 VoD system architecture

The Regional News Service is accessed through a WWW interface, describing the individual news clips that can be accessed. This page includes a brief title and a still image for each clip. When a particular news clip on the Regional News Service page is chosen by the user, the service control replies with a URL that invokes the LAVA video player on the client machine. As soon as the LAVA video player is launched, the indicated video server is accessed and playback of the requested news clip is started automatically. The Regional News Service is described in detail in section 3.

The production process of the News Service is shown in figure 2 and described in more detail in section 4. The local news transmission from Tromsø is recorded locally and digitized in real-time to M-JPEG by a cooperating research institute in this town - NORUT. Immediately after recording these files are transferred via the Supernett to NR in Oslo. At NR the files are transcoded to MPEG [5] during the night and laid out via a Word Wide Web user interface in the morning, ready for viewing over the Internet.

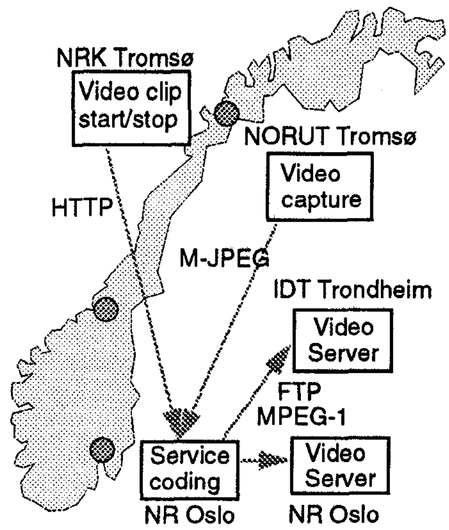

Figure 2 Local news production process

This paper is organised as follows: In section 2, we discuss work related to video on demand systems and digital broadcasting and video player technologies. In section 3, we discuss the requirements for and describe the implementation of the Internet-based version of the Regional News Service we have developed for NRK. In section 4 we describe the production process for this news-on-demand system. In section 5 we describe the LAVA video player implementation, while section 6 concludes the paper.

2 RELATED WORK

We describe other web-based TV services to our knowledge, and briefly discuss the technologies in use. We keep our attention to systems that support real-time playback over network connections, thus excluding file-based video players.

2.1 Related Web-based TV channels

Today, distribution of TV is dominated by the use of commercial products like VDOLive and StreamWorks (described below). Several examples of video demonstration Web sites exist, but distribution of regular TV programs via Internet is still in its infancy.

NBC broadcasts live business and financial news. This is a regular traditional TV broadcast available on Internet. However, previous programs are not stored subsequent for on-demand access. Distribution is based on real-time playback over Internet using StreamWorks from XING Technologies. Transmission bandwidths varies from 12kbps to 120kbps.

CBS News Up-to-the-Minute is a live news broadcast. There is no storage of previous programs. Distribution is handled by XING StreamWorks.

CNN Newsroom is a daily program targeted for classroom education in USA, with topics like geography, cultural reports etc. Internet Newsroom makes this program available over Internet, and further provides users with a library of previous programs [6]. Like the service described in this paper, Internet Newsroom is available on-demand. Also news clips are encoded as MPEG-1 system streams. Unlike LAVA, one has to download video clips for viewing, instead of playing them in real-time over the network. Timely downloads are compensated for by caching video clips closer to recipients.

2.2 Related Video Player Technologies

Internet real-time video playback systems are now commercially available, and the examples above show their use. Some of these systems could function adequately for the NRK regional TV service. However, as outlined in section 3.1, we have extended the system requirements beyond solely playback. The player should be able to play video in real-time from a network connection, use MPEG-based audio/video coding and provide the user with VCR control. When the LAVA project started, no available technology could meet these requirements. Thus, we developed a playback system that would allow us to experiment with several aspects of video on demand technologies. In the meantime, several similar player technologies have emerged and are discussed below. None of these technologies meet all our requirements, but provide examples of the rapid developing field of network based digital video player technology.

StreamWorks is a commercial product form XING Technology. StreamWorks can play audio and video down to 9600 bit/second. Video- and audio quality is limited at low bandwidths. At least an ISDN equivalent connection is needed for high quality audio or video and audio. Compression is MPEG based, and playback is done in a video player client connected to a server over the Internet. StreamWorks does not offer any VCR control.

VDOLive is made by VDOnet Corporation. This player can function as a Netscape plug-in or stand alone. Video is encoded using Wavelet-based compression [7]. This method is claimed to give better compression ratios than MPEG for comparable quality of video. Similar to StreamWorks, VDOLive use a client server architecture and does not provide VCR control.

InterVU Inc. develops a plug-in MPEG player for Netscape called PreVU. It plays MPEG video while caching video to a local disk file. Thereafter, the video can be replayed with maximum quality. In contrast to StreamWorks and VDO-Live, PreVU is not dependent on special server software for playback. It utilizes the HTTP protocol for video transfer. PreVU does not currently handle audio, and there is no VCR control.

The distributed systems group at Oregon Graduate Institute of Science and Technology has made an MPEG player that plays synchronized MPEG video streams combined with Sun μ-law audio over Internet [8]. This system adapts dynamically to changes in the execution environment using software feedback mechanisms and dynamic quality of service control. A client-server relationship is established for playback. This player offers VCR functionality.

3 THE REGIONAL NEWS SERVICE

This section discusses the design requirements for the regional news-on-demand service, and describe the WWW-based implementation of the service.

3.1 Service Requirements

In designing the infrastructure for the regional news service, we have emphasized different aspects, most of them critical to NRK in the future development of their digital video infrastructure. The first requirement is to use standards in the coding of video. Within the LAVA project we also investigate MPEG-2 over ATM networks for NRK [9], but in the public oriented service MPEG-1 was the natural choice due to greater availability in user's equipment and lower bandwidth requirements. We are aware that there are other methods of video encoding that in certain cases may be more efficient than MPEG, but the emphasis on standards is so strong that we wanted to use MPEG as a foundation for the video coding.

The second requirement is that this system should have a link to other activities within NRK. The link to digital broadcasting in the future based on MPEG-2 is obvious, but there is also a link to archiving problems. The current archive within NRK is a text based archive where video is stored separately on cassettes and have to be operated manually to check a reference in the archive. Coding the transmissions in MPEG-1, gives the possibility to develop and experiment with digital video archives where a search based on keywords can immediately lead to a corresponding digital video clip [10].

The third requirement is to ensure a certain size of the user group to the experimental system. We therefore chose UNIX workstations with MPEG-1 software decoding as the initial test platform, since this is the major platform within the Universities of Norway connected by the Supernett. The choice of MPEG-1 also offers a smooth upgrade path from software decoding to hardware-supported decoding, immediately utilizing MPEG chips when these become more widespread on PC's, giving even better video quality.

The fourth requirement is to use video coding that gives a certain bandwidth. The LAVA project is oriented towards broadband networks and it would not be interesting for us to work with bandwidths that are accessible on a modem level. Currently the lowest bandwidth offered is 150 kbit/s. The video is offered at a higher rate for those who have access to more powerful networks and workstations.

3.2 Service Implementation

A screen page of the regional news service is shown in figure 3. Today's news is presented on the entering page and each news clip in the news program can be viewed separately. A clickable image from each section is presented with a still image and a descriptive text. Only material that have passed an intellectual property right evaluation are presented. Due to storage limitations, only news from the last month is accessible from the server.

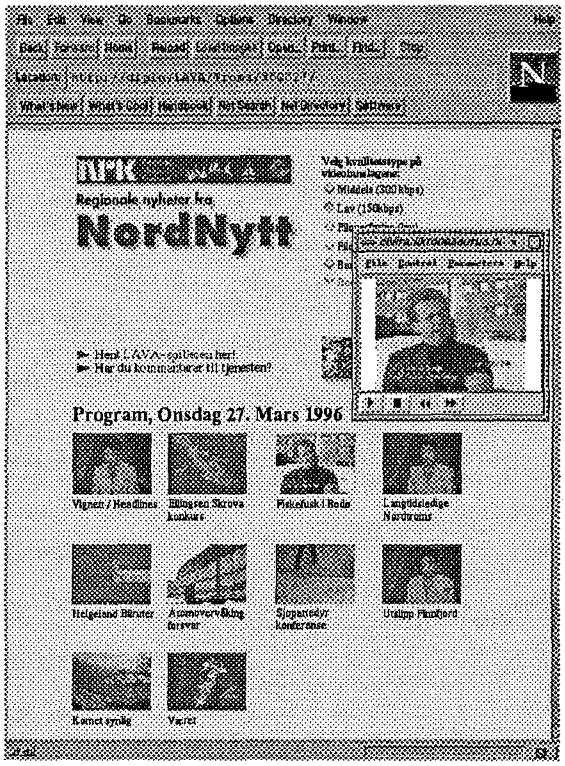

Figure 3 The Regional News Service

To be able to play the video clips on demand, the LAVA player software must be downloaded and installed on the client workstation. The player is described in more detail in section 5.

Two video qualities are provided; 150 kbit/s and 300 kbit/s which corresponds to a resolution of 192 * 144 and 256 * 192. With a frame rate of 25 fps, these are the resolutions that our reference platform (Sun SparcStations and SGI Indy) can manage to decode in real-time with reasonable quality using a software MPEG-1 decoder. Even TV professionals, who are normally very critical to digital video quality, consider this quality to approach acceptable quality for news programs.

In order to reach users who are connected to a lower bandwidth network, video clips can be transferred as a file and played locally. The service also provides the user with access to audio only.

Load balancing between available video servers is done in a round robin fashion and is controlled by the Web-service. In this particular setup with one server at NR and one server at IDT, we alternate between the two available servers.

It is important for NRK to get feedback from the users of the service. For this reason the user must register some demographic data (education, age, gender etc.) and provide a username and password before they can access the service. After using the service for a while, the user will be asked to answer a questionnaire. Typical questions are about the usability of the service, how often such a service will be used, the importance of the demand aspect of the service, ideas for future interesting services etc.

4 DISTRIBUTED PRODUCTION

To produce "Nordnytt" for the Internet involves several steps, including registration of the news clip, capturing, file transfer from capturing location to encoding location, encoding of MPEG, storage using the Elvira video server and make it accessible for the public using WWW.

4.1 Locations involved in the process

The distributed production requires several steps and involves four physical locations in the process:

- NRK produces "Nordnytt" in their regional office in Tromsø. Each day a new newscast is produced and published in this office. The news is broadcast in the evening and lasts for 15 minutes. The program includes several short news clips (from 20 seconds to a few minutes).
- NORUT is a research institute also located in Tromsø, and have been responsible for the real-time video capture of the newscast. The capturing of "Nordnytt" is done by an SGI Indy workstation with a Galileo video board. This board is capable of capturing 25 JPEG frames per second, generating an audio/video stream with bit rate depended on the chosen JPEG quality factor. We use a bit-rate stream of 15 Mbit/s. This indicates a 1.5 GByte file for 15 minutes recording each day.
- NR in Oslo runs several SGI machines with MPEG encoder software. The M-JPEG file is cut in several news clips and are distributed on the local area network for a M-JPEG to MPEG transcoding and finally stored in the local Elvira video server. At the same time, the Regional News Service homepage is updated.
- IDT in Trondheim operates the other Elvira video server and receives MPEG files and directory information once the video clips are encoded at NR.

Details of the production process are discussed in the next section.

4.2 Generating MPEG News Clips

Each night, a new 1.5 Gbyte file arrives at NR, ready for a M-JPEG to MPEG transcoding. Once the file has arrived 3 types of operations are started (see figure 4).

Splitting

Before encoding, the file is cut into several files, each containing a particular clip in the program. The encoding of MPEG is a highly CPU bound process. Generating 300 kbit/s MPEG video on an Indigo2 is a factor of 1:55 (1 seconds of MPEG video is encoded in 55 seconds).

Since the cutting is done on the boundary between two news clips, information about when each clip starts and the length has to be recorded before this process can start. A "script" at

NRK, following the production during the live broadcast, registers these parameters during the program. After the program is finished this person fills in an HTML-fill-out-form to communicate today's program to the split and encoding process. This form contains a sequence of attributes (one such sequence for each news clip):

- **Title.** This field is used as a header for each clip in the WEB-page that presents toady's news. See figure 3.
- **Description.** This fields gives a short description of the clip, not covered by the title and snapshot (optional).
- **Start and Length.** This field indicates where in the original video file this clip is located.
- **Internet.** This field is marked "false" if the news clip contains material where the "intellectual property rights" are not resolved.

MPEG Encoding

Once the file is cut into separate M-JPEG clips the encoding process is initiated. This will eventually produce 2 MPEG system streams for each news clip. Each of these streams address one of the two quality domains (150 and 300 kbit/s).

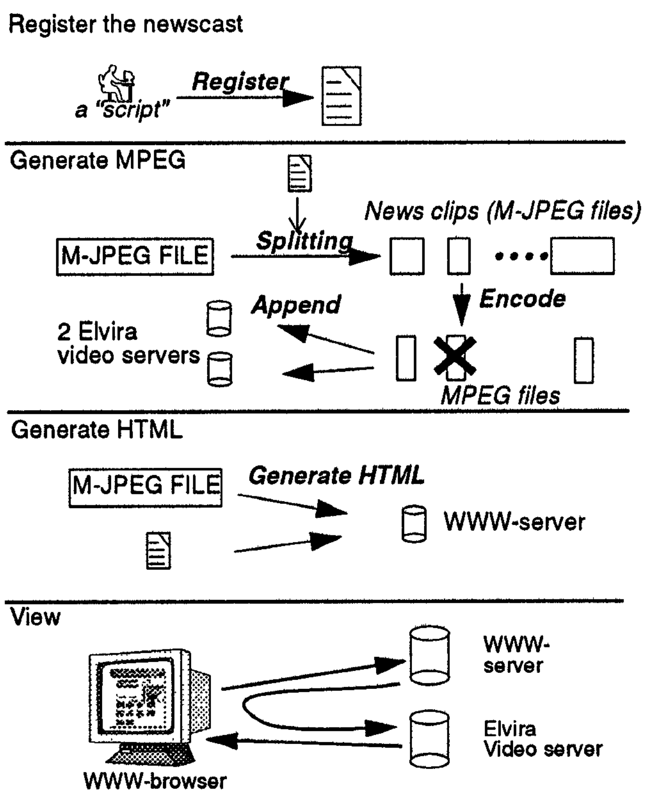

Figure 4 The MPEG encoding and registration

To meet the encoding workload, a coarse grained parallelism is needed. The only possible solution with the current software-based encoding, is to split the video sequence and parallelize the encoding process on available computers with MPEG encoding capability (currently 3 SGI machines).

The complete MPEG-encoding process takes 13-15 hours (includes ftp-time for the M-JPEG file). The production of today's Nordnytt is therefore ready for the Internet community 13-15 hours later than the live TV-broadcast.

The Video Server and WWW

Finally, the MPEG-files have to be registered in the video server and the HTML-page of today has to be included in the WWW server. In order to prevent overloading we have duplicates of the video server. One is located at NR and the other at IDT where the Elvira audio/video server has been implemented [3]. Each MPEG-encoded file is copied into the server file system and is registered in a directory structure. The video clips are now accessible using the LAVA player.

The video server is accessed through a WWW-interface as shown in figure 3. This HTML file describing today's news is produced by use of the title and the optional description field from the registration and a single image from each clip. In addition to today's news, the video server contains news for up to a month. Storage capacity limits long time storage. News clips for one day fills 60 Mbytes of disk space.

5 THE PLAYER IMPLEMENTATION

The LAVA digital video player can read MPEG video from a network connection in a client-server architecture and offer file playback as well. The player offers some VCR capabilities and adjusts the decoding process according to available CPU resources. The implementation is based on the MPEG video decoder developed by the Berkeley Plateau Multimedia Research Group [11][12][13] and the MPEG audio decode library[1] by Tobias Bading, Technical University of Berlin.

5.1 Player Stages

In principle, the LAVA player has five stages as viewed from the main data flow when the player is operating:

- input stage
- demultiplexing stage
- synchronization stage
- decode stage, and
- output stage

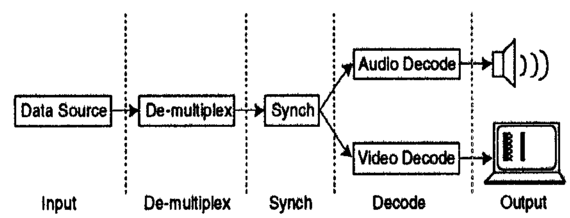

Figure 5 MPEG-1 Player data flow

1. *ftp://ftp.uni-passau.de/pub/unix/sound/maplay/*

Input Stage

The input to the player is either a file or data from a network connection. The input stage provides a bitstream interface to the demultiplex stage. Thus the demultiplex stage can request data from input in the form of one or more bits.

Demultiplex Stage

The demultiplexing stage takes care of splitting an MPEG-1 system stream into MPEG video and audio elementary streams respectively. Here, time stamps associated with video and audio data are extracted and passed to the synchronization stage.

Synchronization Stage

The synchronization stage controls the playing of MPEG-1 system streams. It generates calls to the demultiplex stage, controls time stamps, and decides whether to skip frames or not by setting flags before calling the decoding stage.

Decode Stage

Decoding is separated into a video and audio part. The video part decodes an MPEG video stream output from the demultiplexer stage, putting video frames into a scheduling queue for display by the output stage. The audio part works in a similar manner.

Output Stage

As with decoding, output is also separated into one video and one audio part.

Audio output must handle different audio devices. We have designed a generic interface towards the MPEG decode stage, and several different audio solutions are supported through this interface:

- SGI Audio Library
- SUN μ-law and PCM
- LINUX VoxWare (PC soundcards: e.g. SoundBlaster)

The video part interfaces to a window system that displays decoded video frames in a video window. As of this writing, we only support the X window system.

5.2 Player-Server Communication

The Elvira video server archives the news, which is a set of MPEG files contained in a simple directory structure. The MPEG video clips can be requested using a standard URL syntax, indicating that an Elvira protocol will be used.

```
Elvira://<videoserver>/<videoclip-id>
e.g.
Elvira://kr.nr.no/960418/film1-150kbps.mps
```

Figure 6 Elvira URL format

Once the client has performed the initial access to the server using the Elvira protocol and thus established a control channel, the video server will set up a TCP playback channel back to the machine running the LAVA video player. The control channel remains in place for the remainder of the player session, and is used for our simple VCR protocol between the LAVA client and the Elvira video server.

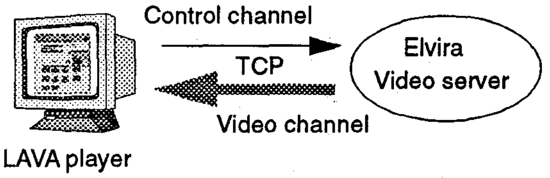

Figure 7 Client-server data flow

The VCR capability of the client is very simple. The client/server protocol is request driven, and the network object in the client issues continuous requests to keep the playback buffer filled at all times. A stop is simply implemented by not sending out further requests, and the video playback is subsequently stopped. Latency for the response of VCR controls is depending on the amount of buffering, which is dynamically changing, but averages to one I-frame. The user currently experience about 0.5 second response time on Play/Stop VCR controls.

In the regional news service described in section 3, we use low bit-rate MPEG streams of a few hundred kilobits/sec. For this application, TCP is a well suited communication protocol because of the inherent flow control and reliability. We need these protocol features, since the Berkeley-based MPEG-1 decoder [12] is very sensitive to packet loss. UDP could have been used if we were very careful about network load to avoid packet loss, but we traded slight jitter introduced by TCP retransmissions in the case of packet loss, for the ease of not implementing an error-tolerant MPEG decoder. We found that TCP did not add substantial overhead in the LAVA player, and decided to postpone introduction of lossy or more advanced protocols with better real-time properties.

LAVA clients are located anywhere in the ATM-based Supernett, using traditional IP over ATM for the client-server connection. Client machines with traditional network connections (e.g. Ethernet) are using the service as well. The software MPEG decoding is CPU intensive, so the service has so far been limited to users with workstation class machines and moderate (> 150 kbit/s) to high capacity (> 10 Mbit/s) network connections.

Since the LAVA MPEG player is implemented in software, speed of the decoding process is what limits our maximum MPEG speed. The decoder implementation is single-threaded and is not optimized for speed. We found that our workstations could support an MPEG stream of maximum 300 kbit/s in order to keep up with the decoding process. We would like to stress the network with high capacity MPEG streams, but this requires hardware decoder implementations. To support higher speed and hardware decoding is a goal of the next generation LAVA client that is discussed in section 6.2.

5.3 Synchronization

When decoding MPEG in software, we have experienced lack of machine resources for doing full frame rate playback even on the most powerful workstations. Since the LAVA player is required to execute on a range of machines, it should adjust to their performance and load.

When full quality decode is impossible, we must select which parts of the MPEG-1 stream should be favoured for decode and presentation, while keeping audio and video synchronized. Several aspects must be considered in this respect:

- Making pictures correspond with audio.
- The frame rate at which the player is capable of decoding video and maintaining full sound quality.
- Various target platforms where the player is supposed to work.
- Workload of client machines.

A common method in handling synchronization is to give priority to audio and skip video when the software decode process cannot keep up with real-time. If we ensure that sufficient audio data is present in the audio hardware buffers at any time, this gives a reproduced audio that will sound natural without drop-outs. We are now left with the task of displaying video frames according to timing data embedded in the MPEG-1 system stream. We extract these time stamps in the demultiplex stage and pass them on with the associated elementary streams to the decode stages.

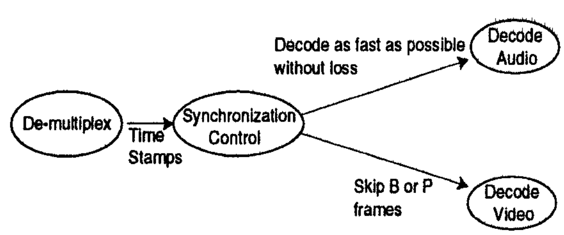

Figure 8 Synchronization mechanism

As figure 8 shows, all time stamps are passed from the demultiplexer to the synchronization control algorithm. This algorithm decides whether frames must be skipped or not, and set appropriate flags to signal this state to the decode stage. As the decode algorithm is about to start, it checks the flags before commencing. The player supports three types of synchronization:

- Frame skipping and clock adjustment
- Buffer management
- Frame rate control

Frame Skipping and Clock Adjustment

The LAVA video player handles all frame skipping on the client side. This wastes bandwidth since data sent from the video server is not used at the client. If the Elvira video server could parse MPEG system streams, frame skipping could be done by the server when necessary [14].

An internal software clock in the LAVA video player registers real-time. This clock is compared to the PTS and the SCR extracted from the MPEG-1 system stream to decide if it should:

- skip frames (or disable frame skipping when we can resume full decoding again).
- readjust internal real-time clock.

If frames must be skipped, we start with B-frames and pass I- and P-frames through for decoding. If this does not bring the video in synchronization with the internal clock, we start skipping P-frames as a last resort.

In order to support the frame skipping synchronization above, we re-adjust the internal clock with timestamps from the MPEG-1 stream by aligning the internal software clock with the last SCR found by the demultiplexer. These values will drift apart due to halts in player execution caused by manual pausing, paging on heavily loaded machines, network delays and other interruptions of the program. This mechanism makes the player robust against all types of execution interruption that it is exposed to in a multitasking environment without hard real-time support.

Buffer Management

Skipping frames results in the video decoder stage requesting more data from the demultiplexer stage. The video decode stage collects data from the elementary video stream bounded buffer between these stages. We must ensure that this buffer holds sufficient data such that the video decoder gets the next frame not to be skipped.

In the worst case, the synchronization control algorithm has decided that we must skip both P- and B-frames. In this case the decoder will scan through the video stream until it finds the next I-frame.

The amount of data in the buffer is controlled using watermark levels as shown in figure 9. When the player control algorithm decides to skip frames (B or P), it raises the low watermark in the video elementary stream buffer. Thereby, the player control algorithm will delay decoding until the data level is above the low watermark, thus first calling the demultiplexer (one or more times). The level is lowered back to its original value when the player resumes normal operation where it doesn't skip frames. The buffer that smooth out playback is always large enough to hold at least one I-frame.

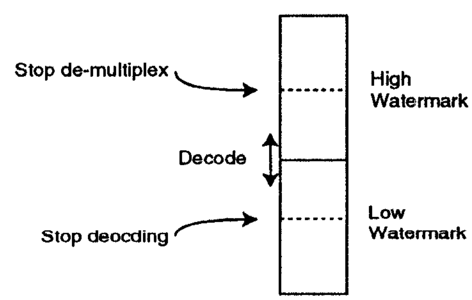

Figure 9 Player data stream buffer

One must also take into account that loading lots of video data into the video elementary stream buffer also effects the capacity necessary in the audio elementary stream buffer. This is due to the fact that audio and video are interleaved in the MPEG-1 system stream. By extracting video, audio will consequently come as a by-product of this process.

Frame Rate Control

The output stage handles presentation of decoded video frames such that they are displayed at the correct time. This time is computed from the frame rate information found in the MPEG-1 stream when the player software initializes.

Decoded video frames are entered into a queue together with their presentation time. The display stage is delayed until the first frame in the queue is due for presentation.

Since we are CPU limited on video decode, we have not been able to work with playback rates that would stress the network severely. This is an area of further work and will be discussed in the next section.

6 CONCLUSION AND FURTHER WORK

Our conclusions fall in two categories. First, we summarize our experience and give directions for future work related to the regional news service and digital broadcasting in general. Second, we discuss ideas for improvement of the LAVA video player.

6.1 Regional News Service

The regional news service has been in trial use from April 1996, and has gained a considerable interest in Norwegian media. We have already registered several hundred users and they are generally pleased with the ability to interactively chose news clips and access the regional newscasts independent of time and location.

This small-scale experiment provides important insight in the usefulness of similar services in a larger scale. Initial response indicates that the video and audio quality provided by this service is adequate for news on demand services. The technology has also been met with considerable interest from professionals developing distance education, electronic documentation, library and museum applications and services.

NRK is currently working to utilize technology from the regional news service in a new program series. This program will have nationwide coverage, and users can use our system to browse through selected parts of previous programs and interactive parts not accessible in the regular broadcast.

The production of the regional news service is based on the available infrastructure at the involved research institutes. If a similar service should be provided for the general public and for a wide range of programs, the broadcast corporation must employ large-volume hardware-based MPEG encoders, and further research must solve a significant storage and server capacity problem.

There has been a significant interest among users of the regional TV service for access from the home over low-speed (modem) connections as well. However, we have chosen to concentrate on high end networks and workstations and believe that general advance in computer architecture and network infrastructure will solve the network capacity and video decoding problem in a relatively short timeframe.

Deregulation of telecommunications in Europe from 1998 will lead to strong competition, more bandwidth and lower prized services. There is also a development towards heterogeneous networks, including the integration of ATM, CATV, ADSL and mobile network technologies that will provide users with new access points and higher bandwidths to the home market.

On the client side, introduction of the new multimedia PC standard (MPC-III) may solve many of the end system problems experienced using current terminal equipment. In particular, low-priced hardware MPEG decoders and high speed network interfaces are expected to be provided in consumer PC systems. Our experiments have been performed in a high speed networking environment, but our findings indicate that similar results can be obtained in general networking environments. However, switched technologies has the advantage of isolating network load to active users only, and our experience from this project indicate that some form of switched network technologies are necessary to support video services of this kind.

Interactive TV services can also be presented on a regular TV – provided that some logic for interactivity is added – when set top boxes based on MPEG video and audio are deployed more broadly. The digital television broadcast system in Europe (DVB) will be based on MPEG, and the introduction of the next generation CD-ROM (DVD) will further strengthen the push towards MPEG in consumer equipment.

While technical problems are being solved, there exists a significant problem related to intellectual property rights. For example, we had to provide a system that could solve the problem that NRK does not always have the rights to redistribute news footage from international sources.

6.2 LAVA MPEG Player

The LAVA MPEG player has proved to be a user friendly helper application easily integrated with a WWW browser and the WWW-based service control. However, during the regional news field trial and our own experimentation we have encountered many ideas for improvement of the player. These ideas are briefly summarized in the following.

MPEG decoding is a highly CPU bound process and is the primary bottleneck in the LAVA player. The decoding process should be done in dedicated hardware, but at this stage we have found that software decoding gives the required level of portability across different client machines. A problem in the decoding software is an inefficient synchronization between audio and video decoding because of the single threaded architecture of the LAVA player. For instance, the video decode module blocks while waiting for a correct display time. This time could for example be used to decode audio.

The VCR control is limited, and we would like to develop a more advanced client-server protocol that could control the transfer rate directly from the video server. So far, the client is the one that discards data that is not needed for display, thus spending more network resources than necessary.

The LAVA project is currently developing a PC-based MPEG player to broaden the user community. The player has an open design with software MPEG decoding that can easily be further developed to support hardware decoders. The LAVA player will be integrated with the DeVise Hypermedia System developed at University of Aarhus, Denmark [15] to support inclusion of hypermedia anchors inside a video stream. We also consider to offer the MPEG player as a Netscape plug-in module.

Use of more robust video encoding techniques like MPEG-2, give us the opportunity to utilize lossy protocols like UDP, direct AAL5 ATM access and real-time protocols like RTP and IPng developed to support continuous media with real-time requirements. The ATM-based high speed network used in LAVA is currently based on TCP/IP over permanent ATM virtual circuits. In addition to new protocols, we would like a dynamically reconfigurable network that allows many high speed users to access the service simultaneously. This is a challenging scalability problem both in terms of video server capacity as well as network configuration and capacity.

7 ACKNOWLEDGEMENTS

The LAVA project has mainly been financed by the Norwegian Research Council with additional support from Uninett AS. The LAVA project team has worked closely with NRK and we are indebted to Hege Skoglund for her advises and encouragement throughout the project. We are also grateful for contributions from Asbjørn Oskal, NORUT, Tromsø and Stein Langørgen, Roger Midtstraum and Olav Sandstå at IDT, Trondheim that have developed the Elvira video server used in this project and assisted in the design of the MPEG player. We are also grateful for contributions from Terje Dalen and Lasse Øverlier that have implemented early versions of the MPEG player. Professor Stein Gjessing at the University of Oslo has provided valuable comments on the manuscript.

8 REFERENCES

[1] Langørgen, S.: *Design and documentation of the Elvira System*, December, 1995, http://www.idt.unit.no/~videodb/elvira/design.ps

[2] Lovett, H.: *Publishing Video over the Net. A presentation of the LAVA project* (Delivery of Video over ATM) Position paper at Ercim WWW Working Group (W4G) conference, Oxford, November 21 1995. *See http://www.cis.rl.ac.uk/events/W4G/papers.html*

[3] Sandstå, O., Midtstraum, R. and Langørgen, S., *Design and Implementation of the Elvira Video Server*, To appear in Proceedings of Norsk Informatikkonferanse 1996, Alta, Norway, November 18-20, 1996.

[4] Hjelsvold, R.*VideoSTAR - A Database for Video Information Sharing*, Dr.Ing. Thesis, Norwegian Institute of Technology, November 1995. See *http://www.idt.unit.no/IDT/grupper/DB-grp/tech_papers/hjelsvold95.html*

[5] Gall, D. *MPEG: a video compression standard for multimedia applications*, Communications of the ACM, 1991.

[6] Compton, C. L and Bosco P. D. *Internet CNN NEWSROOM: A Digital Video News Magazine and Library*, Center for Advanced Engineering Study Electrical Engineering and Computer Science Department, Massachusetts Institute of Technology.

[7] Graps, A. *An Introduction to Wavelets*, IEEE Computational Science and Engineering, Summer 1995, vol.2, num. 2, IEEE Computer Society.

[8] Cen, S. et al. *A Distributed Real-Time MPEG Video Audio Player*, in Proceedings of the Fifth International Workshop on Network and Operating System Support of Digital Audio and Video (NOSSDAV '95), Durham, New Hampshire.

[9] Bryhni, H. *An empirical evaluation of MPEG-2 video transport over an ATM network*, Technical report, IMEDIA0695, Norwegian Computing Center, December 20, 1995. See *http://www.nr.no/home/hilde/LAVA/LAVA_publ.html*

[10] Hjelsvold, R., Langørgen, S., Midtstraum, R. and Sandstå, O. *Integrated Video Archive Tools*, Proceedings of ACM Multimedia'95, San Francisco, CA, November, 1995.

[11] Rowe, L. A. *Continuous Media Applications*, Multipoint Workshop, ACM Multimedia 1994, San Francisco, CA. See also: *http://bmrc.berkeley.edu/projects/mpeg/*

[12] Patel, K., Smith, B.C. and Rowe, L.A. *Performance of a Software MPEG Video Decoder*, Proc. of ACM Multimedia 93, Anaheim, CA, August 1993, pp.75-82.

[13] Rowe, L. A., Patel, K.D, Smith, B.C, Liu, K.: *MPEG Video in Software: Representation, Transmission and Playback*, In Proceedings of the IS&T/SPIE Symposium on Electronic Imaging Science and Technology, San Jose, CA, February, 1994.

[14] Han, C-C and Shin, K. G.: *Scheduling MPEG-Compressed Video Streams with Firm Deadline Constraints*, Proceedings of ACM Multimedia '95, San Francisco, CA., November 1995.

[15] Grønbæk, K. and Trigg, R.H. *Hypermedia System Design Applying the Dexter Model*, Guest editors' introduction to the special issue on Hypermedia in Communications of the ACM 37, 2, 1994, pp. 26-29.