

Muksis project

Application report

Richard Domander
Tuomas Mäenpää
Teemu Nisu
Tommi Teistelä



Version: 0.2
Public
December 8, 2008

University of Jyväskylä
Department of Mathematical Information Technology
Jyväskylä

Approver	Date	Signature	Name verification
Project manager	__.__.2008		
Customer	__.__.2008		
Supervisor	__.__.2008		

Document information

Authors:

- | | | |
|-------------------------|-----------------|-------------|
| • Richard Domander (RD) | dimadoma@jyu.fi | 050-3482668 |
| • Tuomas Mäenpää (TM) | tutamaen@jyu.fi | 040-7600465 |
| • Teemu Nisu (TN) | tejonisu@jyu.fi | 040-8349310 |
| • Tommi Teistelä (TT) | totateis@jyu.fi | 045-6528709 |

Document title: Muksis project, Application report

Pages: 23

Source file: application_report-0.2.tex

Abstract: This is the application report of software project Muksis. This document introduces to the reader the existing architecture and inner workings of MPlayer (mainly version 1.0-rc1), and explores in detail the parts relevant to the features implemented in Muksis project.

Keywords: Software architecture, black frame detection, DVB subtitles, seek in MPEG TS, MPlayer

Contents

1	Introduction	2
2	MPlayer	3
3	Leffakone - the Target Environment	5
3.1	Usage of Black Frame Detection	5
4	Black frame detection	7
4.1	Background	7
4.2	Implementation	7
5	MPEG TS seek	11
5.1	Background	11
5.2	Implementation	11
6	DVB subtitles	13
6.1	Background	13
6.2	Implementation	14
6.3	Version differences	16
7	Audio filtering	17
7.1	Dynamic range compression	17
7.2	Normalization	18
8	File formats and Data structures	20
8.1	EDL file	20
8.2	Structure and contents of MPEG TS	20
9	Analysis	22
10	Testing	23

The following acronyms and terms appear in this document:

BFVF Black Frame Video Filter

CLUT Color look up table

Demultiplexing The act of separating elementary streams from MPEG TS

DVB subs Digital Video Broadcasting Subtitles

EDL Edit Decision List

EOF End of File

Elementary stream A video, audio or subtitle stream carried in MPEG TS

FIFO First In First Out

GOP Group of Pictures

GUI Graphical User Interface

MPEG TS Moving Picture Expert Group Transport Stream

PES Packetized Elementary Stream

PIC Programmable Interface Controller

PID Packet Identifier

PTS Presentation Time Stamp

1 Introduction

The original purpose of the Muksis project was to provide commercial skipping and DVB subtitle support for TV recordings in Matthieu Weber's Leffakone home theater environment. Weber had already implemented these features for an older version of MPlayer (pre4) back in 2005, so we tried to either port or reimplement them for a newer version of MPlayer.

Commercial skipping required black frame detection to find the commercial breaks, and more accurate MPEG TS file seeking. DVB subtitle support required decoding, drawing and timing of subtitle packets. It also required various changes throughout MPlayer, because the program wasn't designed to work with TV recordings, only broadcasts. Weber also came up with an additional requirement: he wanted a possibility to normalize and compress audio during playback.

We set out with the latest stable release of MPlayer at the time (1.0-rc2), but later on Weber noticed he couldn't compile it in Leffakone environment, so we reverted back to version 1.0-rc1. During the project we decided to port the features for the development (or SVN) version of MPlayer too. The differences between versions affected mostly DVB subtitle support.

When we started the project, we thought that the implementation of the features could prove to be a challenge, because the team didn't have much experience in working with a larger application that is written in a non-object-oriented language. We estimated that handling the coherence and integration of the components is the most demanding task, but the components proved to be very independent of each other. The differences between the versions of MPlayer affected mostly the implementation of DVB subtitle support.

Chapter 2 introduces the general architecture of MPlayer, and explain which modules are used with MPEG TS files. Chapter 3 explains how MPlayer, and the features implemented in the project are used in Leffakone environment. Chapter 4 describes black frame detection, 6 DVB subtitle support, 5 seeking, and 7 audio filtering in more detail. Chapter 8 explores the various file formats and data structures used in the project. In chapter 9 the implementation of the features is analyzed, and in 10 their testing is explained. All the chapters apply to MPlayer version 1.0-rc1 unless otherwise specified.

2 MPlayer

MPlayer is a cross-platform, open source media player, whose architecture can be seen as an instance of the "Pipes and Filters" architectural pattern. This follows naturally from the process of media replay on computers. There are myriads of codecs, (de)multiplexers, and video processing filters for myriads of different medias, video, audio and subtitle formats, video output streams, video effects and platforms. The software is highly modular, but the responsibilities of the modules are somewhat ill-defined, and everything revolves around `mplayer.c`. The structure of the program has improved since MPlayer version 1.0-pre4, but is still in need of heavy refactoring.

`Mplayer.c` runs the show on a high level: it handles and delegates user input, keeps the play list up to date, and figures out the necessary components for playback (decoder, demultiplexer, subtitles, media handling...). With command line parameters the user can, amongst other things, choose a video codec, subtitle and audio language, and add various video filters and change their individual parameters. In the MPlayer software package there's also MEncoder which is a versatile video decoding, encoding and filtering tool.

Figure 2.1 explains what happens when you play a MPEG TS file with MPlayer. First the individual streams are demultiplexed in `demuxer_ts.c`. During the replay the user can seek the file (change position in the streams), or supply MPlayer an EDL file which, for example, makes MPlayer skip commercials. All seek commands are passed to `demuxer_ts.c`, because it knows how to change position in the streams, and how to resynchronize them after seeking. Before presentation the streams are decoded, for example, `dvbsub.c` handles the decoding of DVB subtitles. Various effects (audio compression, video de-interlacing...) can be added to video and audio streams before presentation with filters. Black frame detection (`vf_bf.c`) also acts as a filter, but instead of processing the picture data, it only marks down all the black frames it finds into an EDL file.

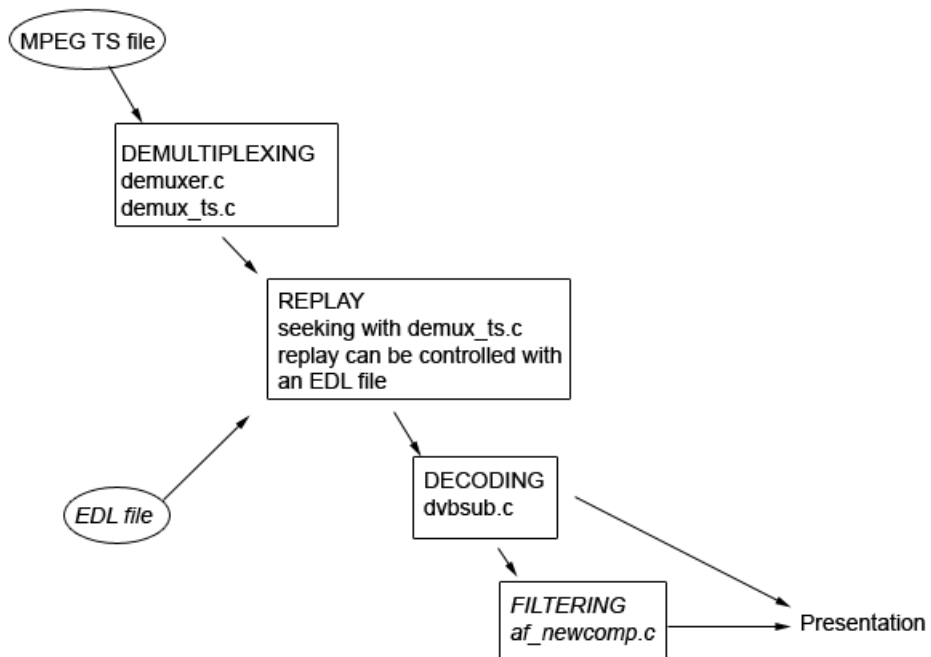


Figure 2.1: General architecture of MPlayer 1.0-rc1

3 Leffakone - the Target Environment

Leffakone is Matthieu Weber's home theater PC. On hardware level it consists of a DVB tuner, a timer device, an infrared receiver, a graphic adapter and a sound card. A TV antenna is connected to the DVB tuner. The timer is used via parallel port, and is connected to the Wake-on-LAN, so that the computer can be powered up when it's time to record a TV program. The infrared sensor is used to pick up signals from a remote control.

On the software level Leffakone consists of the assembly code for the PIC of the timer device, the python code for automating recording and shutting down the computer after recording, Lirc for handling infrared signals, and Freevo, which is a personal video recorder application for Linux. Freevo uses MPlayer (version 1.0-rc1) to display and record TV broadcasts. Freevo uses X Window System for its GUI.

3.1 Usage of Black Frame Detection

In Leffakone environment black frame detection is used in two phases. In phase one (see figure 3.2) there are two instances of MPlayer running simultaneously: one that records the TV broadcast, and one that starts shortly after the other to create an EDL file (see chapter 8) of the recording. The recording is made by "dumping" the streams from the MPEG TS broadcast with `demux_ts.c`. Only one video, audio and subtitle stream is recorded. The instance which creates the EDL file runs the recording through the BFVF. In phase two (see figure 2.1) the recording opened with the newly created EDL file, which then controls the playback of the file, i.e. commercial breaks are skipped.

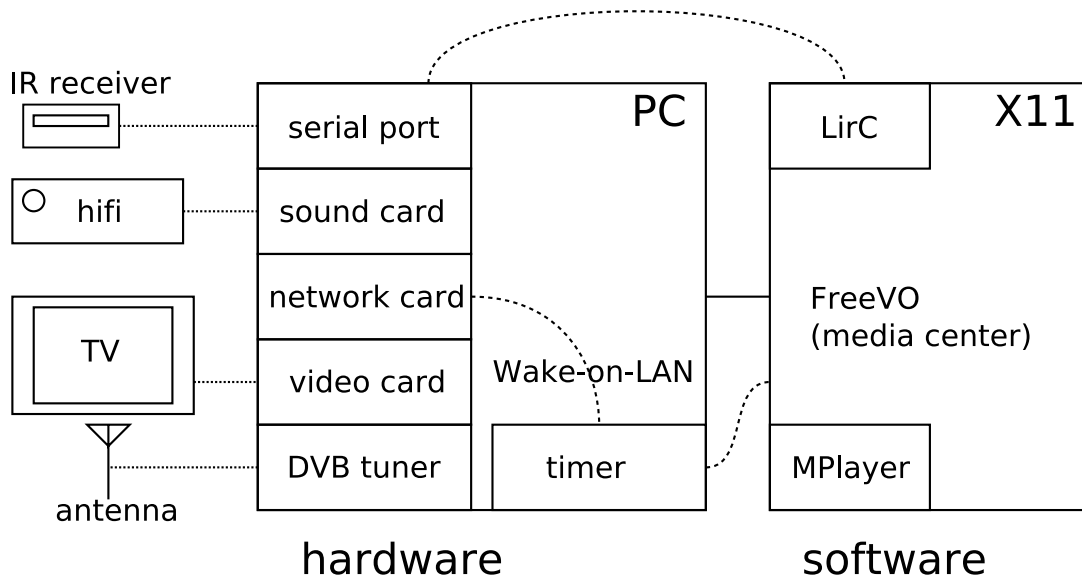


Figure 3.1: The structure of Leffakone

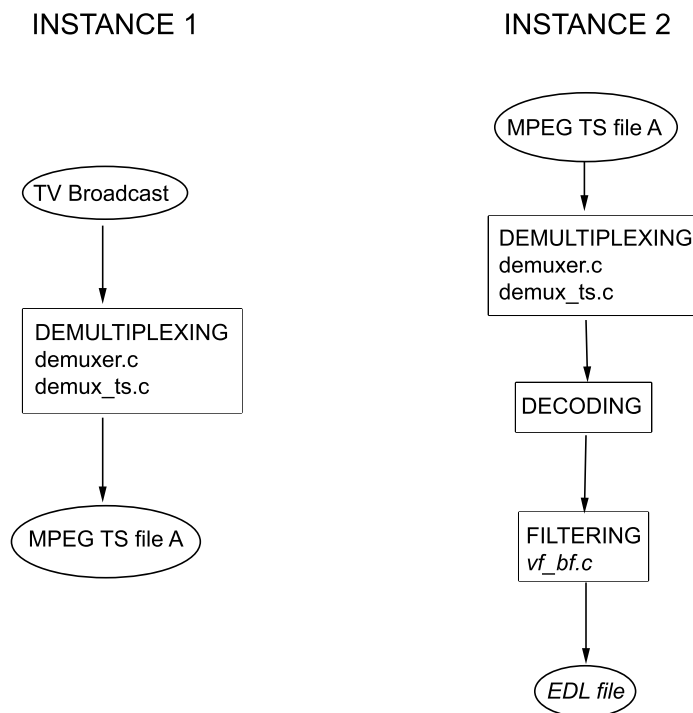


Figure 3.2: Creation of a TV recording and EDL file in Leffakone

4 Black frame detection

4.1 Background

The purpose of black frame detection is to create an EDL file, which can be used to skip commercials during playback of TV recordings. BFVF can be used to detect commercials because commercials are short sequences between black frames. When the user requests video filters with command line parameters, MPlayer passes the request to `vf.c`, which manages the video filter chain. `vf.c` then parses the command line parameters, opens the specified filters, supplies them their individual parameters, and arranges them to a linked list. During playback decoded video frames pass through this chain of filters before they reach video out (see figure 4.1).

Video filters are very independent. In addition to programming the BFVF itself, you only have to mention the new filter in `vf.c` and the `Makefile` of `libmpcodec` so that MPlayer and the compiler find it. While processing the video BFVF's `put_image()` function writes the PTSs of the detected commercials into an EDL file. The file is opened in BFVF's `open()` function, and freed in `uninit()`. Together with accurate MPEG TS seeking the generated EDL file can be used to skip commercials during playback. The EDL file is described in detail in chapter 8. The structure of MPlayer has improved since version `pre4`, and you don't need to get the PTS of a decoded frame as an external variable from `dec_video.c` anymore. Nowadays there's a version of `put_image()`, which has PTS as one of its parameters. The parameter is passed to it by `vf.c`.

4.2 Implementation

Black frame detection was ported from Matthieu Weber's old code. Some refactoring was done, for example `calcing()`, which decides whether a frame is black or not, was cleaned up. There was some redundancy (copy-paste code), "magic numbers" i.e. constants with no explanation, and variables with confusing names. The detection algorithm itself wasn't changed. While doing "printf debugging" we noticed that `vf_config()` gets called every time video parameters change, for example, a commercial may use different aspect ratio than the program, so no initializations should be done in this function (the old code did). There was also a memory




Figure 4.1: Architecture of black frame video filtering

leak: `vf_priv_s` structure was never deallocated, now `free()` for the structure is called in `vf_uninit()`.

`vf_priv_s` structure holds the pointer to the EDL file to be written, and variables necessary for detection of commercials: length of the current sequence between black frames, time when the previous black frame sequence started, length of the previous black frame sequence etc.

`Calcimg()`, which is called in `put_image()`, decides whether a frame is black or not by calculating the average luminance value of the frame, and marking down the highest individual luminance value. If both values are below their corresponding thresholds, the frame is considered to be black.

`Put_image()` is called once for every video frame. It creates the markings for the resulting EDL file. A possible commercial break is considered to begin when a black frame is detected, i.e. a black sequence begins. The length of the non-black sequence following the black sequence is measured. If the length of the non-black sequence before a new black sequence is shorter than `ad_max_len` the sequence is considered to be an ad. As long as the program detects new ads the end of the commercial break

is moved to the middle of the latest black sequence (see figure 4.2). A commercial break ends when a non-black sequence longer than `ad_max_len` is encountered. If EOF is encountered, and the length of the sequence between the latest black frame and EOF is shorter than `ad_max_len` then the end of the commercial break is set to EOF (note: this is done in `vf_uninit()`!).

In Weber's patch the commercial break was marked to the EDL file from the beginning of the starting black sequence to the end of the ending black sequence. We changed it so that it now marks the skip from the middle of the starting black sequence to the middle of the ending black sequence. This was changed because of the MPEG format, i.e. if we jump to a frame that is not the key frame of the GOP an ugly "mosaic" is seen. This doesn't bother as much if the frame is black. Also, the short black sequence now seen in the middle of the program lets the user know a jump occurred. We think it's now less confusing if a new program starts after a commercial break.

The thresholds, EDL filename and `ad_max_len` can be changed by the user with command line parameters. The default values for the thresholds and `ad_max_len` were empirically determined by Weber, and he claims the program attains 99% accuracy with them. The parameters are parsed in `vf_open()`. We changed the implementation so that an error message is printed if all the given parameters cannot be parsed. Default value is used for the faulty parameter, and all the parameters following it. Default values are also used if the user gives no parameters.

A mention of our BFVF was added to `vf.c` and Makefile. In versions `pre4`, `rc1`, and `rc2` the Makefile in sub folder `libmpcodecs` needs to be changed. In the current development version of MPlayer you modify the Makefile in the root folder of the program.



Figure 4.2: Commercial detection in BFVF

5 MPEG TS seek

5.1 Background

In order to skip commercials during playback, we need to be able to seek accurately, i.e. to change position in the streams in MPEG TS. During playback `mplayer.c` waits for keyboard events, and stores the detected key presses into a FIFO buffer. If a seek command (i.e. a key mapped to seek is pressed) is detected, a seek is initiated by `mplayer.c`. A seek can also be automatically initiated by an EDL file. First `mplayer.c` calls the seek function of the generic demultiplexer (`demuxer.c`). The generic demultiplexer checks if the file format and selected media allow seeking (for example, DVD allows, TV does not), calculates the distance of the seek target from the current file position, and makes other necessary preparations (resets the synchronization of the demultiplexed streams, for example). The generic demultiplexer has a function pointer to the seek function of the format or media specific demultiplexer (in this case `demux_ts.c`). The specific seek function does the actual seeking. After the seek in `demux_ts.c` is done, that is: after the correct new position in the video is found and audio is synchronized with video, `mplayer.c` calls `update_subs()` function in `mpcommon.c` (see chapter 6). `update_subs()` initiates the subtitle format specific tasks needed to display current subtitles.

5.2 Implementation

The improved seek code in `demux_ts.c` was redone to study different ways of implementing it. The basic principle remains the same, however.

As the MPEG-TS stream by itself doesn't contain a typical header-like structure to refer to, the seek function must work by estimating the bitrate of the stream and using the MPEG timestamps in it as references. The modified seek function stores the current MPEG video timestamp from the demuxer's state in a temporary variable, seeks in the file stream using the demuxer's estimate of the bitrate and attempts to re-synchronize the video and audio streams by skipping frames until the two match using code similar to the original MPlayer seek function.

Since seeking backwards accurately is generally more difficult, the modified seek function attempts to avoid doing so by performing shorter seeks when moving for-

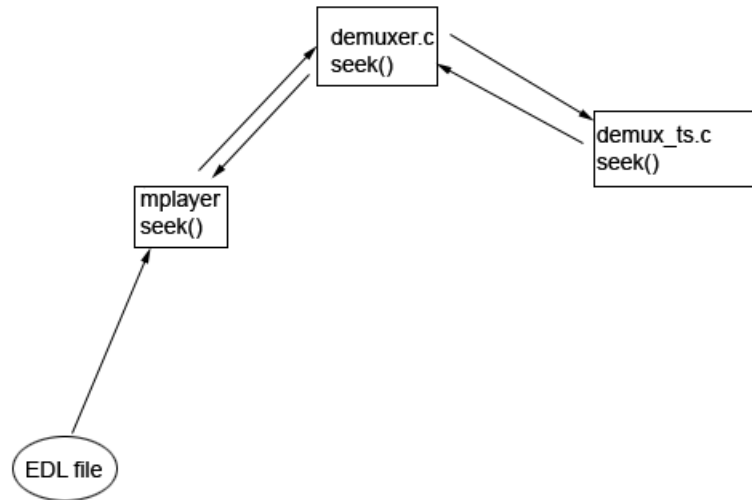


Figure 5.1: Architecture of MPEG TS seeking

ward in the stream, and longer seeks when moving backwards. This also largely avoids the issue of running into the end-of-file when seeking is performed near the EoF and the demuxer has overestimated the stream's bitrate.

After the synchronization the current timestamp in the stream is compared to the one stored earlier. If the difference is more than 0.5 seconds and the seek attempt has been repeated less than 16 times, the seek function attempts to perform another seek using the difference as the seek duration. If the difference is less than 0.5 seconds, the seek function simply returns (seeking more accurately would require implementing the seek functionality on a lower level, taking the different types of MPEG frames into account and possibly requiring modifications elsewhere, too) The possibility of the timestamp values resetting somewhere between the beginning and the end of the seek is dealt with by testing if the new timestamp value is off by more than several minutes in the "wrong direction" (compared to the seek duration) and changing the initial timestamp variable to match the new situation.

6 DVB subtitles

6.1 Background

YLE channels in Finland use DVB subtitles, which are separate from the video stream. In order to show these subtitles, we need to implement timing and decoding of the DVB subtitles. Because both DVD and DVB subtitles are bitmaps (in most cases anyhow) and both formats have support for multiple languages, you can take advantage of the existing code for DVD subtitles in MPlayer. Essentially the data structure `spudec_handle_t` in `spudec_struct.h` (see figure 6.1) acts as a wrapper for DVB subtitles. In our implementation, we utilize the packet queue in this structure, and in the existing code decoded bitmap subtitles are passed on from this structure to the displaying functions.

When MPlayer opens, it initializes DVB subtitles in `mplayer.c` by calling `dvbsub_init_data()`. During playback, and after seeking, `mplayer.c` calls the function `update_dvdsubs()`. `update_dvdsubs()` calls the functions necessary to handle DVD and DVB subtitles. The function deduces the type of the subtitles from the type of the demultiplexer. In case of DVB subs, `update_dvdsubs()` first calls `dvbsub_heartbeat()` in `dvbsub.c`, which checks whether the first packet in the decoding queue should be processed, i.e. is the current time stamp \geq packet time stamp. When a packet is processed in `dvbsub_decode()` the resulting image is handed to the `spudec` structure, from which the image goes on to the displaying functions. After calling `dvbsub_heartbeat()` `update_subs()` loops as long as it can get PES packet data from the buffer of the subtitle stream, i.e. as long as it can call `ds_get_packet_sub()` in `demux.c` successfully. After every call to `ds_get_packet_sub()` `update_subs()` calls `dvbsub_assemble()`, and passes it the received packet data. `Dvbsub_assemble()` adds the packet data to the decoding queue of the `spudec` struct. When a new PES packet is added to the decoding queue in `dvbsub_assemble()`, time stamp and time out are added to these packets to ensure correct decoding, and thus presentation order, time and duration. In summary, `update_dvdsubs()` does two things: it decodes the subtitle packets, and makes sure that they are presented at the right time.

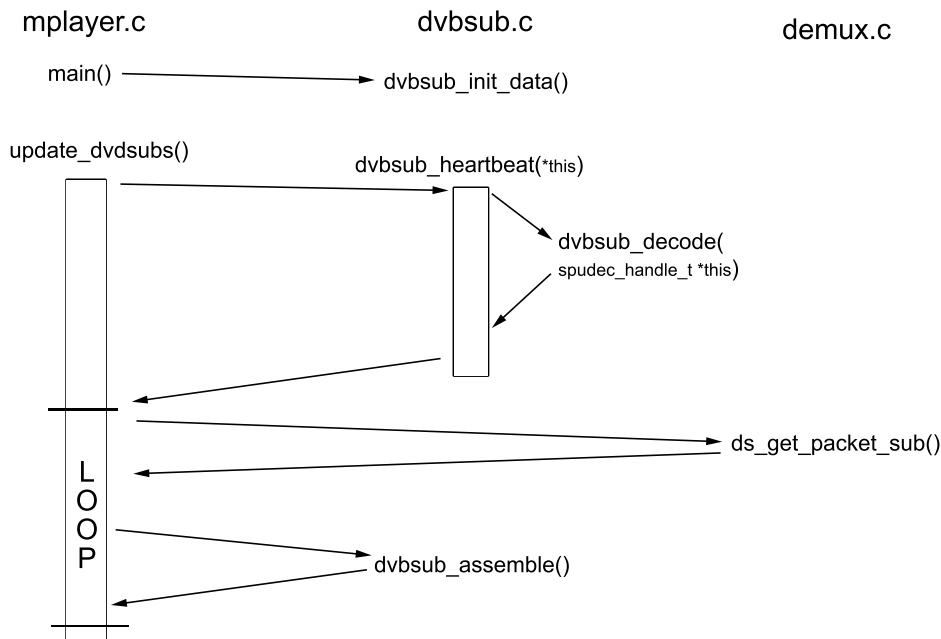


Figure 6.1: Architecture of DVB subtitles

6.2 Implementation

The DVB subtitle packet processing code from Weber’s patch was used almost as is. Only the memory behavior was changed at first, and the drawing of the decoded image was put into one function instead of two. In the patch code memory for the decoded image was deallocated, and then allocated again every time the subtitles changed. Also, the need for memory was dynamically calculated based on the dimensions of the subtitle regions. The code was changed so that the maximum possible amount of memory needed (less than a megabyte) is allocated before the first subtitle packet is decoded. The subtitle regions are erased (filled with 0) before a packet is decoded. This change in memory usage was later removed, because it created a conflict with `vf_scale.c`, and the gain in performance wasn’t worth the time that would’ve been needed to resolve the issue.

During the project Weber noticed that he had supplied us with an outdated version of his code. He then sent us the latest version he had coded. The newer version rendered the subtitles better, so we included the changes it introduced to our version. The old version didn’t initialize the pixel code (the CLUT index of the pixel’s color)

to zero in `decode_4bit_pixel_code_string()`, and it miscalculated the alpha values in `set_palette()`. Because of these errors, the background of the subtitles wasn't completely transparent, and the glyphs didn't have sharp black edges.

During testing we noticed that with some files, the subtitles didn't get shown at all. There were two reasons for this problem. When a MPEG TS file is opened, `ts_detect_streams()` in `demux_ts.c` reads a couple of seconds of the file to find out which elementary streams it contains. The first reason for the problem was that this probing for streams ended after a video and audio stream had been found, i.e. it could end before the probing limit was reached, thus missing possible subtitle streams starting in the beginning of the file. The premature exit condition was changed so that the probing can end before the limit only if also a subtitle stream is found, or no subtitles were requested (`dvdsb_id=-2`). The second reason for the problem was that `demux_ts.c` was designed to work with a TV broadcast, not a TV recording, i.e. `ts_parse()` didn't recognize the type of new elementary streams without the meta data present in a TV broadcast (MPlayer records only one video, audio and subtitle stream). We modified the function `ts_parse()` so that now if it can't recognize the elementary stream by any other means, it reads a packet (with the help of `pes_parse2()`) from the file, and tries to get the type of the stream from it. If the type is not recognized, the packet is discarded. If the type is recognized (to be DVB subtitles) the position in the file is moved back to the beginning of the packet so that the actual contents can be read. Also, the demultiplexer is updated to use this stream, for example if the new stream is a DVB subtitle stream, the subtitle id is changed to PID of the newfound stream. `Ts_parse()` is a rather confusing function with a dual role: it is used both for the (initial) detection of streams, and packet parsing during playback (reading a packet from the file, and then put it to the corresponding output buffer [video, audio or subtitle] of the demultiplexer).

Some changes in `mplayer.c` were needed to make the subtitles work. Several variables in MPlayer refer to the state of different kinds of subtitles. Relevant for DVB subtitles are: `dvdsb->id`, which is used to request bitmap subtitles from the demultiplexer, `sub_auto`, which tells should subtitles be used at all, and `demuxer->sub->id`, which tells what is the PID of the subtitle stream currently used. If `sub_auto=0` (the user used parameter `-noautosub`) DVB subtitles aren't initialized. During playback we check if a subtitle stream was found by comparing `dvbsub->id` to `demuxer->sub->id`. If the variables differ, a subtitle stream has been found, and we set `dvbsub->id` to `demuxer->sub->id`, and call `update_dvdsb()` to show the first packet detected. The functionality now in `update_dvdsb()` was

initially in the `main()` function in `mplayer.c`.

6.3 Version differences

The implementation of DVB subtitles differs a bit between the versions. In version newer than rc2 function `update_subtitles()` in `mpcommon.c` is called instead of `update_dvdsubs()` in `mplayer.c` as is done in rc1.

Instead of `dvbsub->id` and `demuxer->sub->id` as in rc1, `mpctx->d_sub->id` and `mpctx->demuxer->sub->id` are used in rc2 and newer versions to check whether we have found a subtitle stream during playback

In versions newer than rc2 the `spudec` structure is initialized in function `init_vo_spudec()` in `mplayer.c`. In rc1 it is initialized in the `main()` function of `mplayer.c`.

7 Audio filtering

Both the actual "peak" volume level and perceived volume (affected by the frequency distribution and the dynamic range of the sound) tend to vary between different TV broadcasts, possibly more so than between different DVD movies, which MPlayer was originally made to play. MPlayer's current source tree does contain some audio filters intended to control these aspects of the sound being played, but both their documentation and implementation are somewhat lacking for the purposes of Leffakone. The project team has thus decided to look at either improving the existing filters or writing suitable ones from scratch.

MPlayer's audio filter interface resembles that of the video filters: a filter has a standard interface structure containing function pointers for key operations, such as (un)initializing the filter, configuring it and processing data through the filter. This structure is added to the list of audio filters in `af.c` and the rest is handled by MPlayer's existing code. Through this interface, audio filters may receive parameters from the command line or from other parts of the application, such as any of the GUI frontends for MPlayer. Audio filters may be "chained" in the same way as video filters.

7.1 Dynamic range compression

The "dynamic range" of any given information is the range of possible values between its lowest and highest levels. For audio, this can be considered the range between the most quiet (but still relevant for the experience) parts of the sound, such as quiet speech, and the loudest parts of the sound, like gunshots or explosions. The desirability of a high dynamic range in audio mostly depends on the quality of the listening environment; with low-quality speakers or a noisy environment the most subtle parts of high-dynamic range audio will be lost. Because of this the dynamic range of audio is often reduced through dynamic range compression.

The Muksis project has implemented a new simple "hard-knee" dynamic compression filter (see fig. 7.1) similar to that used by many hardware DVD players when outputting sound to a device expected to be of low quality, such as the built-in speakers of a TV. The filter's source code can be found in `af_newcomp.c`, and it can be controlled from the command line with parameters like `mplayer -af`




Figure 7.1: A "hard-knee" audio compression filter

$\text{newcomp}=0.7:0.5$, where 0.7 is the compression threshold and 0.5 the ratio by which anything exceeding the threshold is compressed.

The exact function used when the input exceeds the filter's threshold is:

$$\text{output} = (\text{input} - \text{threshold}) * \text{scale} + \text{threshold} \quad (7.1)$$

which scales the part exceeding the threshold by the compression filter's scale setting.

7.2 Normalization

The actual "peak" volume level of audio tends to vary significantly from the maximum allowed by its format; this can make for an irritating listening experience. The volume level can be corrected through normalization: applying a multiplier to the audio sample values bringing the peak volume to the maximum level, or close to it – generally, most implementations do not raise the peak volume to the absolute maximum as this can result in distortion due to limitations in various parts of the audio pipeline.

For short audio clips such as recorded music, this can be achieved easily by simply scanning through the data and finding the highest absolute value, then amplifying the sound by the maximum value divided by the highest value found, scaled to whatever peak volume is desired. For very long audio clips or ones being currently streamed, this approach can't work and the normalization filter must instead react to changes in the volume dynamically. The Muksis project has been looking at ways to do this in MPlayer.

MPlayer's own volume normalization filter in `af_volnorm.c`, while apparently quite old and unmaintained in the `rc1` version, was found adequate for the task.

The filter can take two parameters, which are the compression method to be used (the filter implements two, for some reason) and the average level the normalization filter tries to maintain, given in a syntax similar to the other audio filters. When run with no parameters, the filter tries to keep the peak volume at the maximum level allowed.

8 File formats and Data structures

8.1 EDL file

In our implementation EDL files are used to skip commercials during playback. These files are automatically created by BFVF (see chapter 4). The format is very simple, you only need begin and end timestamps (PTS) of the actions, and the types of the actions, i.e. do you want to mute (=1) or skip (=0) a part of the video. Below is an example EDL file.

Example EDL file

Begin	End	Type
8372.18	8409.26	0
8595.30	8869.18	0
9724.62	9965.10	0
10855.34	11084.90	0
11814.46	12008.66	0

8.2 Structure and contents of MPEG TS

In Finland digital television is broadcasted in MPEG TS format. The broadcast consists of several MPEG TSs, called "channel packets". One MPEG TS consists of one or more multiplexed programs, which consist of several streams. These streams consist of PES packets, which can contain video, audio, subtitle, meta, or other, so called private data (teletext, network information, encryption key for scrambled channels...). These packets carry PTSs, which are used for timing. The meta data tells which streams belong to which program. In TV recordings made by Leffakone (see chapter 3) there's only one video, audio and subtitle stream present. During playback the demultiplexer separates these streams from each other.

The video format used in MPEG TS is MPEG 2. MPEG 2 videos consists of groups of pictures (GOP), which hold 0.5 seconds of video. GOPs consist of different kinds of frames: I-frames, P-frames and B-frames. I-frames are the key frames, they are full pictures. P-frames only carry data which is used to create a displayed frame by modifying the previous I-frame. B-frames are more complicated bi-predictive




Figure 8.1: Structure of the MPEG Transport Stream

frames, i.e. they consist of modifications created by analyzing both the frame before and after the B-frame.

A DVB subtitle packet holds one page instance, which is an unchanging set of subtitles. A page instance consists of several regions (areas upon which subtitles can appear), which consist of objects. One object depicts one glyph. A page instance can also be an "end of display set", which is used to signal that the current subtitles can be cleared.

9 Analysis

10 Testing