

The Problem

- You host a video on your web app
- You want high quality so you host a large 1080p mp4
 - The entire file is 100's of MB
- Every user visiting your page has to download the entire video before playback can begin
 - Very slow to load
 - Entire file must download even for users who will only watch for a few seconds

Chunking

- Avoid the requirement of downloading the entire video before it plays • Provided a way to request short segments of the video Download one "chunk" of the video at a time.
- Typically ~2-10 seconds of playback
- Advantages:
 - Only a few seconds need to be downloaded before playback starts If the user skips around in the video, only request they chunk they
 - skipped to
 - If the user leaves the page without finishing the video, the entire file doesn't have to be downloaded

Chunking - Range

- Even mp4 videos can be chunked
- header
- The response code should be 206 Partial content
- The response should contain a Content-Range header
 - file
 - eg. Content-Range: 10000-20000/500000

• The browser may end a request for a large file that includes a Range

• The Range header specifies a range of bytes being requested eg. A request for "/video.mp4" with a header "Range: 10000-20000" is request the 9999th through 19999th bytes of video.mp4 • Contains the range of bytes being sent and the to total size of the

• eg. Content-Range: 0-499999/500000 if the entire file is sent

Chunking - Range

- Using the Range header can be effective in some cases, however:
 - It adds extra complexity to the server • It relies on the browser to request useful ranges Some browsers might not implement Range and ask

 - for the entire file

We would like a more robust solution

range of bytes..

- User only watches a few seconds -> Only need to request the first few segments User skips to the middle of the video -> Request the
- We host a single video in multiple files • Each file contains a few seconds of playback • The browser requests each segment as needed
- - middle segment

Chunking - Multiple Files

Instead of relying on the browser asking for a specific

HLS vs MPEG-DASH

- segments/chunks: HLS and MPEG-DASH
- HTTP Live Streaming (HLS)
 - Developed by Apple
 - Only supports the H.264 encoding for video
 - Wide-spread adaptation
 - Spec freely available in RFC8216
- - Developed by Moving Picture Experts Group (MPEG)
 - Supports any video encodings
 - No support on Apple devices

• Two major protocols support the idea of breaking a video into smaller

• Dynamic Adaptive Streaming over HTTP (MPEG-DASH)

Spec published as ISO/IEC 23009-1:2022 - Available for \$245 (!)

٦II	space.m3u
∠ TS	space0.ts
⊿ TS	space1.ts
∠ TS	space2.ts
⊿ TS	space3.ts
⊿ TS	space4.ts
⊿ TS	space5.ts
⊿ TS	space6.ts
⊿ TS	space7.ts
⊿ TS	space8.ts
⊿ TS	space9.ts
⊿ TS	space10.ts
⊿ TS	space11.ts
⊿ TS	space12.ts
⊿ TS	space13.ts

#EXTM3U
#EXT-X-VERSION:3
#EXT-X-TARGETDURATION
#EXT-X-MEDIA-SEQUENCE
#EXTINF:6.773433,
space0.ts
#EXTINF:8.341667,
space1.ts
#EXTINF:8.341667,
space2.ts
#EXTINF:8.341667,
space3.ts
#EXTINF:8.341667,
space4.ts
#EXTINF:8.341667,
space5.ts
#EXTINF:8.341667,
space6.ts
#EXTINF:8.341667,
space7.ts
#EXTINF:8.341667,
space8.ts
#EXTINF:8.341667,
space9.ts
#EXTINF:8.341667,
space10.ts
#EXTINF:8.341667,
space11.ts
#EXTINF:6.973633,
space12.ts
#EXTINF:2.836167,
space13.ts
#EXT-X-ENDLIST

:8

HLS

- Divide the video into multiple .ts files
 MPEG Transport Stream files
- One .m3u8 index file containing information about each .ts file and how they combine into a single video
- Your server hosts all files
- Set the video source as the index file
- Browser reads the index file to know when to request each ts file



HLS - Transcoding

ffmpeg -i space.mp4 -hls_list_size 0 -f hls space.m3u8

- Use ffmpeg to convert to HLS
- "-f hls" to specify the output format as HLS
- "-hls_list_size 0" to keep all ts files in the index
 - By deafult, ffmpeg will only keep the last 5 ts files in the index file
 - This is good if you are live-streaming (This is the HTTP) Live Streaming protocol after all)
 - Since our use case is hosting Video on Demand (VOD), we want to keep every ts file in the index
 - Setting the list size to 0 means the size is not limited

chunk-stream0-00001.m4s	5
chunk-stream0-00002.m4s	7
chunk-stream0-00003.m4s	8
chunk-stream0-00004.m4s	9
chunk-stream0-00005.m4s	11
chunk-stream0-00006.m4s	12
chunk-stream0-00007.m4s	13
chunk-stream0-00008.m4s	14
chunk-stream0-00009.m4s	16
chunk-stream0-00010.m4s	17
chunk-stream0-00011.m4s	18
chunk-stream0-00012.m4s	19
chunk-stream0-00013.m4s	21
chunk-stream0-00014.m4s	22
init-stream0.m4s	23
space.mpd	24
	23

28

29

xml version="1.0" encoding="utf-8"?
<mpd td="" urn:mpeg:dash:schema:mpd:2011"<="" xmlns:xsi="http://www.w3.org/2001/XMLS</th></tr><tr><td>xmlns="></mpd>
xmlns:xlink="http://www.w3.org/1999/xli
xsi:schemaLocation="urn:mpeg:DASH:schem
profiles="urn:mpeg:dash:profile:isoff-l
type="static"
mediaPresentationDuration="PT1M48.3S"
maxSegmentDuration="PT5.0S"
minBufferTime="PT13.9S">
<programinformation></programinformation>
<servicedescription id="0"></servicedescription>
<period id="0" start="PT0.0S"></period>
<adaptationset <="" contenttype='</th></tr><tr><th><pre><Representation id="0" mimeType</pre></th></tr><tr><th><SegmentTemplate timescale=</th></tr><tr><th><SegmentTimeline></th></tr><tr><th><S t="0" d="203203' id="0" th=""></adaptationset>
<s 209209"="" d="250250" r="10</th></tr><tr><th><S d="></s>
<s d="85085"></s>
<pre></pre>
<pre></pre>
<pre></pre>

Schema-instance"

.nk" Live:2011"

video" startWithSAI e="video/mp4" codec: "30000" initializa[.]

NPEG-DASH

- Divide the video into multiple files that can use a variety of formats (m4s by default in ffmpeg output which is mp4 encoded)
- One .mpd (Media Presentation) Description) index file containing tons of information in an XML format
- Hosts all files and set the video source as the index file
- Browser reads the index file to find an init file and naming convention for content files to request







MPEG-DASH

1	xml version="1.0" encoding="utf-8"?
2	<mpd <="" th="" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"></mpd>
3	xmlns="urn:mpeg:dash:schema:mpd:2011"
4	xmlns:xlink="http://www.w3.org/1999/xlink"
5	xsi:schemaLocation="urn:mpeg:DASH:schema:MPD:2011 http://standards.iso.org/ittf/Pub
6	profiles="urn:mpeg:dash:profile:isoff-live:2011"
7	type="static"
8	mediaPresentationDuration="PT1M48.3S"
9	maxSegmentDuration="PT5.0S"
10	minBufferTime="PT13.9S">
11	<programinformation></programinformation>
12	
13	<servicedescription id="0"></servicedescription>
14	
15	<period id="0" start="PT0.0S"></period>
16	<pre><adaptationset 0"="" 30000"="" bandwidth="</pre></th></tr><tr><th>18</th><th><pre><SegmentTemplate timescale=" codecs="avc1.640028" contenttype="video" d="203203" id="0" initialization="init-stream\$Represent</pre></th></tr><tr><th>19</th><th><pre><SegmentTimeline></pre></th></tr><tr><th>20</th><th><S t=" mimetype="video/mp4" segmentalignment="true</pre></th></tr><tr><th>17</th><th><pre><Representation id=" startwithsap="1"></adaptationset></pre>
21	<s d="250250" r="10"></s>
22	<s d="209209"></s>
23	<s d="85085"></s>
24	
25	
26	
27	
28	
29	

• Full mpd file from the example

liclyAvailableStandards/MPEG-DASH_schema_files/DASH-MPD.xsd"

ue" bitstreamSwitching="true" frameRate="30000/1001" maxWidth="1920" maxHeight="1080" par="16:9" lang="eng"> "2403076" width="1920" height="1080" sar="1:1">

tationID\$.m4s" media="chunk-stream\$RepresentationID\$-\$Number%05d\$.m4s" startNumber="1">



MPEG-DASH - Transcoding

ffmpeg —i space.mp4 —f dash space.mpd

- Use ffmpeg to convert to MPEG-DASH
- "-f dash" to specify the output format as MPEG-DASH
- Creates an mpd with the name you provide
 - All other files follow a default naming convention
 - May want to create a new directory for each video to avoid naming conflicts

Adaptive Bit-Rate Streaming

The Problem

- You host a video on your web app
 - 2-10 seconds chunks
- You want high quality so you host chunks in 4K@60Hz
 - Can require ~25Mb/s bandwidth to stream
- - The video buffers, stutters, or doesn't play at all
- We need a solution that:

 - Delivers high quality to users with high-speed Internet

• You even use HLS or MPEG-DASH to segment the video into

• And someone visits your site using eduroam on a bad day...

Allows users with slow connections to enjoy your content

- - resolutions
- User visits your page
 - Their browser adapts to the current download bandwidth available

 Instead of hosting the a single video at a single resolution Host multiple versions of the same video at different

• Each resolution requires a different bit rate to stream

Stream the highest bit rate video that fits the bandwidth

• Using HLS or MPEG-DASH

- Create chunks at several different resolutions/bit rates Add information about all resolutions in the index file
- With the video segmented into ~2-10 second chucks Easy for the browser to switch between resolutions

• The browser can adapt the requested bit-rate based on current conditions • Limited interruption for the user, though quality can change over time



wifi slows down

wifi speeds up

Adaptive Bit-Rate - HLS

٦	main.m3u8
٦llı	media_0.m3u8
٦llı	media_1.m3u8
٦III	media_2.m3u8

- Using HLS, m3u8 index files can be nested
- Convert your video into multiple HLS resolutions
- Combine them into a single index file with references to the others
- different resolutions, and 1 audio index file

1	#EXTM3U
2	#EXT-X-VERSION:7
3	#EXT-X-MEDIA:TYPE=AUDI0,GROUP-ID="group_A1",N
4	#EXT-X-STREAM-INF:BANDWIDTH=131049,RESOLUTION
5	media_0.m3u8
6	
7	#EXT-X-STREAM-INF:BANDWIDTH=1131049,RESOLUTIC
8	media_2.m3u8

• This example contains references to 2 video index files at

NAME="audio_1",DEFAULT=YES,URI="media_1.m3u8" =540x960,CODECS="avc1.64001f,mp4a.40.2",AUDIO="group_A1"

)N=322x572,CODECS="avc1.64001e,mp4a.40.2",AUDIO="group_A1"



- The mpd file can contain multiple resolutions
- The media is represented in multiple layers
 - Period
 - Adaptation Set
 - Representation

15	¢	<period id="0" start="PT0.0S"></period>
16	¢	<pre><adaptationset contenttype="video" id="0" segmentalignment="</pre" startwithsap="1"></adaptationset></pre>
17	e de la constante de la consta	<pre><representation bandwid<="" codecs="avc1.64001f" id="0" mimetype="video/mp4" pre=""></representation></pre>
25	e de la constante de la consta	<pre><representation bandwid<="" codecs="avc1.64001f" id="2" mimetype="video/mp4" pre=""></representation></pre>
33	白	
34		<pre><adaptationset contenttype="audio" id="1" segmentalignment="</pre" startwithsap="1"></adaptationset></pre>
35	e de la constante de la consta	<pre><representation bandwidth<="" codecs="mp4a.40.2" id="1" mimetype="audio/mp4" pre=""></representation></pre>
45	白	
46	白	

in multiple resolutions ed in multiple layers

"true" bitstreamSwitching="true" frameRate="30/1" maxWidth="640" maxHeight="1138" par="9:16" lang="und"> |th="480684" width="540" height="960" sar="1:1"...> |th="1000000" width="640" height="1138" sar="5121:5120"...>



Period
Allows the division of a media file into multiple parts
ie. Chapters
For our purposes, we only need 1 period

15	¢	<pre><period id="0" start="PT0.0S"></period></pre>
16	¢	<adaptationset contenttype="video" id="0" segmentalignment="</td" startwithsap="1"></adaptationset>
17	¢	<pre><representation bandwid<="" codecs="avc1.64001f" id="0" mimetype="video/mp4" pre=""></representation></pre>
25	Ē.	<pre><representation bandwid<="" codecs="avc1.64001f" id="2" mimetype="video/mp4" pre=""></representation></pre>
33	ф	
34	↓	<adaptationset contenttype="audio" id="1" segmentalignment="</td" startwithsap="1"></adaptationset>
35	Ę	<pre><representation bandwidth<="" codecs="mp4a.40.2" id="1" mimetype="audio/mp4" pre=""></representation></pre>
45	ф	
46	ф	

"true" bitstreamSwitching="true" frameRate="30/1" maxWidth="640" maxHeight="1138" par="9:16" lang="und"> th="480684" width="540" height="960" sar="1:1"...> th="1000000" width="640" height="1138" sar="5121:5120"...>



Adaptation Set

- Represents different parts of the media that will all be combined into one experience
- Can contain multiple video/audio adaptation sets
 - eg. One for video, one for audio, and another for closed captions

15	$\overline{\Box}$	<period id="0" start="PT0.0S"></period>
16	¢	<pre><adaptationset contenttype="video" id="0" segmentalignment="</pre" startwithsap="1"></adaptationset></pre>
17	ф –	<pre><representation bandwid<="" codecs="avc1.64001f" id="0" mimetype="video/mp4" pre=""></representation></pre>
25	ф –	<pre><representation bandwid<="" codecs="avc1.64001f" id="2" mimetype="video/mp4" pre=""></representation></pre>
33	白	
34		<adaptationset contenttype="audio" id="1" segmentalignment="</td" startwithsap="1"></adaptationset>
35	ф –	<pre><representation bandwidth<="" codecs="mp4a.40.2" id="1" mimetype="audio/mp4" pre=""></representation></pre>
45	ф	
46	ф –	

The example below has one for video and one for audio

true" bitstreamSwitching="true" frameRate="30/1" maxWidth="640" maxHeight="1138" par="9:16" lang="und"> |th="480684" width="540" height="960" sar="1:1"...> |th="1000000" width="640" height="1138" sar="5121:5120"...>



Representation

- media

þ	<period id="0" start="PT0.0S"></period>
¢	<adaptationset contenttype="video" id="0" segmentalignment="</td" startwithsap="1"></adaptationset>
₽	<pre><representation bandwid<="" codecs="avc1.64001f" id="0" mimetype="video/mp4" pre=""></representation></pre>
₽	<pre><representation bandwid<="" codecs="avc1.64001f" id="2" mimetype="video/mp4" pre=""></representation></pre>
ф	
¢	<adaptationset contenttype="audio" id="1" segmentalignment="</td" startwithsap="1"></adaptationset>
₽	<pre><representation bandwidth<="" codecs="mp4a.40.2" id="1" mimetype="audio/mp4" pre=""></representation></pre>
φ.	
白	

 Each adaptation set can have multiple representations • This is where we add multiple different bit-rates of the

 The browser will choose one representation for each adaptation set at a given time based on bandwidth

> true" bitstreamSwitching="true" frameRate="30/1" maxWidth="640" maxHeight="1138" par="9:16" lang="und"> |th="480684" width="540" height="960" sar="1:1"...> th="1000000" width="640" height="1138" sar="5121:5120"...>



Representation

- can choose
 - 480kb of 1Mb bandwidth
- track
 - Only 128kb bandwidth

15	¢	<period id="0" start="PT0.0S"></period>
16	¢	<pre><adaptationset contenttype="video" id="0" segmentalignment="</pre" startwithsap="1"></adaptationset></pre>
17	_ ₽	<pre><representation bandwid<="" codecs="avc1.64001f" id="0" mimetype="video/mp4" pre=""></representation></pre>
25	_ ₽	<representation bandwid<="" codecs="avc1.64001f" id="2" mimetype="video/mp4" td=""></representation>
33	_ ↓	
34	ģ	<pre><adaptationset contenttype="audio" id="1" segmentalignment="</pre" startwithsap="1"></adaptationset></pre>
35	₫	<representation bandwidth<="" codecs="mp4a.40.2" id="1" mimetype="audio/mp4" td=""></representation>
45	_ ¢	
46	þ	

This example has 2 resolutions for video that the browser

• This example only has one representation of the audio



true" bitstreamSwitching="true" frameRate="30/1" maxWidth="640" maxHeight="1138" par="9:16" lang="und"> |th="480684" width="540" height="960" sar="1:1"...> |th="1000000" width="640" height="1138" sar="5121:5120"...>



					4	

chunk-stream0-00001.m4s chunk-stream0-00002.m4s chunk-stream1-00001.m4s chunk-stream1-00002.m4s chunk-stream1-00003.m4s chunk-stream1-00004.m4s chunk-stream2-00001.m4s chunk-stream2-00002.m4s init-stream0.m4s init-stream1.m4s init-stream2.m4s output.mpd

- eg. "chunk-stream0-0000X.m4s" and "init-stream0.m4s" are the files containing the video representation with id == 0
- The browser dynamically requests the specific files desired given the circumstances
- Your server hosts all files

15	¢	<period id="0" start="PT0.0S"></period>
16	¢	<adaptationset contenttype="video" id="0" segmentalignment="</td" startwithsap="1"></adaptationset>
17	_ É	<pre><representation bandwid<="" codecs="avc1.64001f" id="0" mimetype="video/mp4" pre=""></representation></pre>
25	Ē.	<pre><representation bandwid<="" codecs="avc1.64001f" id="2" mimetype="video/mp4" pre=""></representation></pre>
33	ф	
34	¢	<adaptationset contenttype="audio" id="1" segmentalignment="</td" startwithsap="1"></adaptationset>
35	Ē.	<pre><representation bandwidth<="" codecs="mp4a.40.2" id="1" mimetype="audio/mp4" pre=""></representation></pre>
45	φ	
46	ф	

Each representation has it's own set of content files

true" bitstreamSwitching="true" frameRate="30/1" maxWidth="640" maxHeight="1138" par="9:16" lang="und"> |th="480684" width="540" height="960" sar="1:1"...> th="1000000" width="640" height="1138" sar="5121:5120"...>



- Your task [For AO2]:
 - Choose between HLS and MPEG-DASH
 - Programmatically convert an uploaded mp4 into the format you choose with multiple resolutions
 - Serve those videos by setting the source of a video to the index file for that video
- Hint: Get ffmpeg to do all the work for you You should not manually write any of these files. Do research to find the command(s)/flags you need to send
- - to ffmpeg to do the job

- DASH
 - DASH)
- We must use a 3rd party video player
 - Several players available (eg. dash.js)

Video Players

Most built-in video players do not support HLS or MPEG-

• You cannot rely on the browser having a player for either of these formats (eg. During grading we will not use a browser with built-in support for HLS or MPEG-

Examples in the following slides will use video.js

 This example downloads the css and js for video.js from a CDN Uses "class" and "data-setup" attributes on a video element to tell the

library to do its thing

You now have a video player that supports both HLS and MPEG-DASH

1	html
2	<pre> description of the second seco</pre>
3	designed → head>
4	k href="https://vjs.zencdn.net/8.10.0/video-js.c
5	<title>CSE312 Video Example</title>
6	
7	<mark>⇔<body></body></mark>
8	
9	<pre>video class="video-js" width="300" controls autoplay da</pre>
10	<source src="main.m3u8"/>
11	Your browser does not support video playback
12	
13	
14	<pre><script <="" pre="" src="https://vjs.zencdn.net/8.10.0/video.min.js"></script></pre>

Video Players

rel="stylesheet"/>

ta-setup="{}">

</script>

<video class="video-js" width="300" controls autoplay data-setup="{}"> <source src="output.mpd"/>

Your browser does not support video playback </video>



Video Players

- Your video player will now have a consistent look across all browsers
- Don't have to worry about what formats each browser supports
 - You support any format supported by video.js

1	html
2	<mark> </mark> <html lang="en"></html>
3	description of the second s
4	k href="https://vjs.zencdn.net/8.10.0/video-js.cs"
5	<title>CSE312 Video Example</title>
6	└ <mark>│</mark>
7	<mark>⇔<body></body></mark>
8	
9	<video autoplay="" class="video-js" controls="" da<="" p="" width="300"></video>
10	<source src="main.m3u8"/>
11	Your browser does not support video playback
12	└ <mark>/video></mark>
13	
14	<pre><script ;<="" pre="" src="https://vjs.zencdn.net/8.10.0/video.min.js"></script></pre>

rel="stylesheet"/>

ta-setup="{}">

</script>





Live Streaming

Live Streaming

 We've talked about uploading and host mp4 videos using a streaming protocol • A VOD service

• What about live streaming?

 Most live streaming isn't truly live • There will be several seconds of delay in the stream Acceptable loss to gain accuracy

Live Streaming Typical setup (eg. Twitch/YouTube Live/etc.)

- Real-Time Messaging Protocol (RTMP)

 - this case

• The server transcodes the video into a streaming format (eg. HLS/MPEG-DASH) and continually updates/ generates index files

• User streams their video into an ingest server using the

• RTMP is a container for any real-time communication

The content of RTMP happens to be a media stream in

Live Streaming

• When a viewer visits a live stream

- requesting content
- index file
- Repeat until the stream ends

 When a viewer visits the VOD of a past live-stream Serve an index file for then entire stream No different than watching the stream live

The browsers asks for the latest index file and starts

• When it nears the end of that index file, request a new

Live Streaming

- some time:
 - The stream is not truly live

- If the delay is unacceptable (eg. Zoom):
 - Use UDP instead of TCP
 - Do not transcode
 - Accept dropped packets as a part of life

Since the transcoding process of the ingest server takes

 The streamed content is downloaded via TCP/HTTP • Reliable. You will not miss a second of video