Savannah Campbell

December 16, 2016

Digital Preservation

Nicole Martin

The Great Video Codec Debate: JPEG 2000 and FFV1

"I think that too much ink is being spilled over the issue of JPEG 2000 vs FFV1"
– George Blood[1]

As analog video formats become increasingly obsolete and magnetic media degrades over time, video tapes will need to undergo digitization so that their content can be preserved. Through this process, the images and sounds encapsulated in the tape's magnetic particles will be transformed into a series of 0s and 1s: a representation of the video stream that a computer can read and play back. When a cultural heritage institution decides to digitize its video collection for preservation purposes, it becomes necessary for the institution to consider which format should be used for the resulting files. While an uncompressed format would capture all of the data in the video as is, uncompressed encoding results in large files that rapidly fill precious storage space on hard drives and servers. On the other hand, lossy compressed formats like Apple ProRes make digital video files easy to work with in production environments due to their efficient encoding and smaller file sizes, but they compress the image irreversibly—a loss of information that conflicts with the goal of preserving the original content as it was. There is another video codec option for audiovisual archivists: mathematically lossless. With lossless video codecs, the resulting files are still compressed, but in a way that preserves all of the

[1] George Blood. Personal interview by Savannah Campbell. 28 Nov. 2016.

information of the original stream, and the files themselves can be decoded to form the full, original image.

Two mathematically lossless video codecs have become prominent in the field of audiovisual archiving: JPEG 2000 and FFV1. A debate has emerged within the archiving community over which of these two video codecs should be the preferred preservation file format for digitized analog video. No consensus has yet been reached as different archival communities and user groups support each format for various reasons. Taking a look at the histories of JPEG 2000 and FFV1, examining the benefits and drawbacks of each, and investigating which institutions implement them in their video preservation workflows will shed some light on this fierce divide, and why one lossless codec may be preferred over the other.

A Tale of Two Codecs: Their Histories and Users

Developed by the Joint Photographic Expert Group, part one of the JPEG 2000 standard, which defines the format's codestream and compression algorithm, became an International Standard (ISO/IEC 15444-1) in December 2000. Since then, seven additional parts of the standard have been released that, among other things, define different extensions, specify encryption protocols, and add support for three-dimensional images.[2] The file format is scalable and institutions can utilize some or all parts of the standard to their choosing. The standard also specifies that JPEG 2000 can be used for either lossy or lossless compression. The lossless variant of JPEG 2000 has gained support in cultural heritage institutions in large part because the Library of Congress's recommended video preservation format is the lossless JPEG 2000 codec in an MXF wrapper. According to the Library of Congress, JPEG 2000 is "The first practical

[2] "Overview of JPEG 2000," Jpeg.org. Web. <https://jpeg.org/jpeg2000/>.

moving image compression standard that doesn't throw away picture content in any way."[3] With the endorsement from the Library of Congress, many institutions, particularly in the broadcasting and production sectors, have adopted JPEG 2000 in MXF as their chosen preservation format.

It is mostly proprietary software and hardware that support lossless JPEG 2000. For example, Front Porch Media's SAMMA encoders, developed by Jim Lindner, have been used in large-scale digitization projects, such as those performed by the Library of Congress, or others carried out by vendors. As most of the software that supports JPEG 2000 is proprietary, the codec is most commonly used in larger institutions and more broadcast and production-based environments. One thing that users must be wary of, however, is that many applications that claim to offer JPEG 2000 support, such as Avid Media Composer, only work with the lossy variant of the codec, not the lossless one.[4] One common argument for using JPEG 2000 for audiovisual preservation is that it is already used in Digital Cinema Packages, or DCPs, for motion picture distribution. However, DCPs only use lossy JPEG 2000, not lossless JPEG 2000.[5] While not necessarily a deterrent to choosing JPEG 2000 as one's preservation file format, users should be aware of that not all JPEG 2000 files are compressed equally, and this effects not only the quality of their encoded video streams, but what software they are able to work with.

FFV1, short for "FF Video Codec 1," was developed by Michael Niedermayer and first released in 2003. Niedermayer was also a key developer of FFmpeg, an open source command-line tool for encoding, transcoding, and decoding a variety of video and audio formats. FFV1

---

[3] James Snyder, "JPEG 2000 in Moving Image Archiving," 13 May 2011. <http://www.digitizationguidelines.gov/still-image/documents/Snyder.pdf>

[4] Peter Bubestinger, Herman Lewetz, and Marion Jaks, "Comparing Video Codecs and Containers for Archives," 13 Aug. 2015. < http://download.das-werkstatt.com/pb/mthk/info/video/comparison_video_codecs_containers.html>

[5] Peter Bubestinger, "Re: Thoughts on video codecs?" Message to Savannah Campbell, 28 Nov. 2016, Email.

was designed to be integrated into FFmpeg and is part of its codec library.[6] As such, FFV1 is also supported by other applications that are built on or work with FFmpeg, such as the video player VLC and the digital preservation system Archivematica. The specifications for the codec are still being expanded on, with the latest version, FFV1 version 3 (FFV1.3), being released in 2013. FFV1 is still in the process of standardization. Currently, the CELLAR (Codec Engineering for LossLess Archiving and Realtime transmission) working group is developing specifications for FFV1, as well as the container Matroska, and is taking steps to standardize these formats through the Internet Engineering Task Force (IETF). When used together, wrapping a FFV1 video stream within Matroska would constitute a completely open source file format.[7] Though FFV1 can be paired with a number of different container formats, it has predominantly been paired with Matroska in archival contexts, though some archivists such as Peter Bubestinger at the Austrian Mediathek opt to use FFV1 in an AVI wrapper instead.

Though it is still not particularly widely used, FFV1 has seen increasing adoption within archival institutions, particularly ones that do not want to rely on proprietary applications to encode and decode their files. As FFmpeg and FFmpeg-based applications are already a large part of workflows in many audiovisual archives, the built-in support for FFV1 within these applications is desirable, particularly within among groups of archivists that prefer using free and open source tools and formats.

---

[6] "Sustainability of Digital Formats Planning for the Library of Congress: FF Video Codec 1," *digitalpreservation.gov*, 23 Dec. 2015. <http://www.digitalpreservation.gov/formats/fdd/fdd000341.shtml>

[7] Ashley Blewer and Dave Rice, "Status of CELLAR: Update from an IETF Working Group for Matroska and FFV1," *Mediaarea.net*, 03 Oct. 2016. <https://mediaarea.net/Events/PDF/2016-10-03_iPRES_Status_of_CELLAR.pdf>

Lossless Compression Simplified: How JPEG 2000 and FFV1 work

As both JPEG 2000 and FFV1 are mathematically lossless codecs, they share some similarities in how they work to compress video streams. Both formats use intraframe compression. This sets them apart from lossy formats, most of which use interframe compression. With interframe codecs, the compression is applied over a group frames, while the intraframe compression used by FFV1 and JPEG 2000 applies the compression to each individual frame, one at a time.[8] The compression ratios of the two formats are very similar, and in terms of the quality of their compression, FFV1 and JPEG 2000 produce comparable video files.[9]

Though FFV1 and JPEG 2000 both apply compression frame by frame, they use different algorithms to compress the individual frames of their respective bitstreams. JPEG 2000 is compressed in planes, whereas FFV1 uses slices. JPEG 2000 splits the frame into square tiles and treats each block like its own images. Then, JPEG 2000 uses a wavelet transform to compress the matrix of tiles. During this process, as many coefficients as there are pixels in the image are applied one-dimensionally across the matrix of tiles twice: once vertically and once horizontally to constitute the full two-dimensional image. After wavelet transformation, quantization occurs, which simplifies the coefficients. In order for this process to be lossless, a 5/3 wavelet transform must be used, so the quantized values produced are integers that do not need to be rounded. Finally, the frame is then encoded in three separate passes, with the purpose

[8] "Sustainability of Digital Formats Planning for the Library of Congress: FF Video Codec 1."
[9] Bubestinger, et al, "Comparing Video Codecs and Containers for Archives."

of the final pass being to catch any bits that were missed during the first two encoding passes. In this way, JPEG 2000 is designed to minimize encoding errors.[10]

FFV1 is compressed and encoded differently than JPEG 2000. FFV1 is compressed in slices which are scanned and encoded from left to right and line by line, in a similar fashion that the scanlines in an analog video frame are read.[11] Each slice is encoded separately, and then they are merged back together to form each frame. FFV1 also uses a process called multithreading, which can distribute the process of encoding slices across a computer's multiple processors, which allows FFV1 to perform faster than other lossless codecs.[12] Additionally, slices are a variable that can be adjusted. The higher the slice count, the longer the encoding will take, but a higher slice level, such as 24 or 30, is recommended for preservation purposes. This is because FFV1 has built in fixity features. Checksums can be created beyond the frame level, and each slice in each frame can have a checksum attached to it. Slices thus not only reduce any visual artifacts from data corruption to a smaller area of the image, but facilitate pinpointing the exact location of damage or bit corruption.[13]

The differences in how JPEG 2000 and FFV1 work to encode video also impact the speed at which they are able to do so. One thing that archivists on both sides of the aisle seem to agree on is that in terms of efficiency, FFV1 is simply faster at encoding and decoding than JPEG 2000. One of the things that impacted audiovisual archivist Dave Rice's decision to use

---

[10] Jason Elzinga and Keith Feenstra, "JPEG 2000: The Next Compression Standard Using Wavelet Technology," 04 Dec. 2001. <http://faculty.gvsu.edu/aboufade/web/wavelets/student_work/EF/how-works.html>

[11] Michael Nidermayer, "FFV1 Video Codec Specification," *ffmpeg.org*. <http://www.ffmpeg.org/~michael/ffv1.html>

[12] Herman Lewetz, "The MediaConch Presentation at IASA," *mediaarea.net*. 22 Oct. 2015. < https://mediaarea.net/MediaConch/2015/10/22/iasa-presentation/>

[13] Blewer and Rice

FFV1 was that he found, through practice, that JPEG 2000 was very slow by comparison.[14] Even Jim Lindner, who was instrumental in helping the Library of Congress decide on JPEG 2000 as its recommended preservation format, concedes the point that JPEG 2000 is a "CPU hog."[15] When running tests on a variety of lossless codecs, Peter Bubestinger, an archivist at the Austrian Mediathek, found the difference in encoding and decoding speeds between FFV1 and JPEG 2000 to be quite stark. The results of his tests can be seen in the image below.

Video source file:	VQEG reference video "football" (NTSC-SD, 720x486px, 30fps, yuv422p, 8bpc)					
Codec	**Encoding**	**Decoding**	**Filesize**	**% of uncompressed**	**Implementation**	**Details**
Dirac	23 fps	29 fps	122 MiB	50.6%	libschroedinger	log, framemd5
FFV1 (version 1)	55 fps	285 fps	109 MiB	45.2%	libavcodec (FFmpeg)	log, framemd5
FFV1 (version 3)	216 fps	277 fps	111 MiB	46.1%	libavcodec (FFmpeg)	log, framemd5
H.264 lossless	94 fps	190 fps	118 MiB	49.0%	libx264	log, framemd5
JPEG2000 lossless	9.9 fps	51 fps	113 MiB	46.9%	libopenjpeg	log, framemd5

Video source file:	SVT reference video "park joy" (full-HD/1080p, 1920x1080px, 50fps, yuv420p, 8bpc)					
Codec	**Encoding**	**Decoding**	**Filesize**	**% of uncompressed**	**Implementation**	**Details**
Dirac	4.6 fps	5.3 fps	942 MiB	61.3%	libschroedinger	log, framemd5
FFV1 (version 1)	11 fps	64 fps	874 MiB	56.9%	libavcodec (FFmpeg)	log, framemd5
FFV1 (version 3)	31 fps	63 fps	879 MiB	57.2%	libavcodec (FFmpeg)	log, framemd5
H.264 lossless	15 fps	31 fps	957 MiB	62.3%	libx264	log, framemd5
JPEG2000 lossless	1.8 fps	9.3 fps	888 MiB	57.8%	libopenjpeg	log, framemd5

[16]

As seen in the top table, when performing tests with a standard definition file, FFV1 version 3 could be encoded at 216 frames per second and decoded at 277 frames per second.

---

[14] Dave Rice, "Re: Thoughts on video codecs?" Message to Savannah Campbell, 27 Nov. 2016, Email.
[15] Jim Lindner, "Re: Thoughts on video codecs?" Message to Savannah Campbell, 26 Nov. 2016, Email.

[16] Bubestinger, et al, "Comparing Video Codecs and Containers for Archives."

Meanwhile, JPEG 2000 could only encode the same standard definition video at 9.9 frames per second and decode it at 51 frames per second. In terms of how they handle high definition video, the results are similar, with the bottom table again showing that of all of the lossless codecs tested, FFV1 had the fasted encoding and decoding rates while JPEG 2000 had the slowest. Meanwhile, these tests show that FFV1 version 3 and JPEG 2000 maintain a similar compression ratio and overall file size. With slow encoding times the process of creating files increases, and slower decoding speeds hinder playback and quality control checking of files. If an archive has an enormous collection of audiovisual material to be digitized and wants it done quickly and efficiently, FFV1 would appear to get the job done in a timely fashion while maintaining quality. Though speed is not the only factor to consider when selecting a format, FFV1 appears to have the edge over JPEG 2000 in terms of efficiency.

Another thing to consider when choosing a video codec for preservation is what kinds of metadata it supports and how self-descriptive the files are. JPEG 2000 on its own is not very self-descriptive and relies on its container for some of its technical metadata, such as color space information. This is particularly an issue when using JPEG 2000 in an MXF wrapper, as MXF has a very unusual way of storing color space metadata. According to Dave Rice, "The MP4 family, like MOV, JP2, Motion JPEG 2000, and others use the 'colr' atom to store this data. MXF is entirely different from these containers architecturally and doesn't utilize the 'colr' atom." This can cause issues with some decoders when playing back YUV videos, as JPEG 2000-MXF will by default assume the file is RGB. Rice goes on to note that using JPEG 2000 in a different wrapper such as QuickTime may offer better color space support.[17] Meanwhile, FFV1 is highly

[17] Dave Rice, "Re: FFV1 vs other formats for preservation," *Archivematica Google Group*, 24 Sept. 2012.
<https://groups.google.com/forum/#!topic/archivematica/HulV96gJ0go>

self-descriptive and will carry much of its own technical metadata with it on its own, regardless of what container it is put in.[18] Matroska also is highly self-descriptive, but it is not without its faults. For example, Matroska currently does not support timecode tracks.[19] Both JPEG 2000-MXF and FFV1-Matroska have limitations in terms of what technical metadata they support, but FFV1 alone is significantly more self-descriptive than JPEG 2000 on its own.

While error concealment, encoding and decoding speeds, and self-description are all factors to consider when choosing a codec for preservation masters, the primary concern is that these files are accurate representations of the original analog material. JPEG 2000 and FFV1 are both supposed to be mathematically lossless and able to retain all of the original video information. Some archivists on both sides of the debate question the ability of these codecs to do so.

Mathematically Lossless: Too Good to be True?

One of the concerns surrounding the debate among archivists over which video codec should be the preferred preservation format deals with the issue of whether or not these codecs are truly mathematically lossless. Supporters of both FFV1 and JPEG 2000 advocate their chosen format and doubt whether the competing format can be decompressed to reform the full video stream. Jim Lindner, who has been a strong advocate for JPEG 2000 since the Library of Congress was still considering codec options, claims he has "… not seen a single independent engineering test that proves that FFV1 is a completely lossless bitmap…Before JPEG 2000 was accepted there was all sorts of testing done, and I don't see that here."[20] Meanwhile, Peter

[18] Blewer and Rice.
[19] Kieran O'Leary, "Re: Thoughts on video codecs?" Message to Savannah Campbell, 28 Nov. 2016, Email.

[20] Lindner

Bubestinger has worked with both codecs and tested them both rigorously, which ultimately led to his institution implementing FFV1 as their preservation format. According to Bubestinger, "Testing J2K/MXF files and different J2K implementations is why we dropped considering J2K at the Mediathek: Apart from other problems with the files, we couldn't verify losslessness throughout a transcoding workflow."[21] Through tests performed at the Mediathek, Bubestinger was able to reproduce the uncompressed videostream from FFV1 files, but not JPEG 2000 files.

Though advocates of FFV1 doubt the losslessness of JPEG 2000 and vice versa, an evaluation of both formats for their suitability as preservation masters performed by FADGI, the Federal Archival Digitization Guidelines Initiative, seemed to show that both formats were indeed truly mathematically lossless. According to Carl Fleischhauer, a Digital Initiatives Project Manager at the Library of Congress: "As far as I can tell, the two encodings work equally well: after you decode the compressed bitstream, you get back exactly what you started with...I wondered if we were hearing from two different (albeit overlapping communities), each with its own ethos."[22] Though the Library of Congress endorses JPEG 2000 and uses it for their preservation masters, Fleischhauer's assessment of FADGI's findings gives the impression that the two formats are comparable by the Library's standards. This brings into focus the fact that differing archival ideologies, rather than the actual technical proficiencies of each codec, are the cause of the rift between JPEG 2000 and FFV1 adopters. This difference in ethos is particularly reflected in comments on the issue of having standardized formats.

---

[21] Bubestinger, "Re: Thoughts on video codecs?"
[22] Carl Fleischhauer, "Comparing Formats for Digitization," *loc.gov*. 03 Dec. 2014. <http://blogs.loc.gov/thesignal/2014/12/comparing-formats-for-video-digitization/>

The Question of Standardization: How Much Does it Matter?

Moving away from the technical merits and drawbacks of the two formats, the divide among audiovisual archivists over JPEG 2000 and FFV1 appears to be guided by other, more ideological forces. Much of the feelings expressed on this matter revolve around how important it is for a format to be formally standardized. JPEG 2000 has been an ISO standard since 2000, while FFV1 is currently undergoing the standardization process through the IETF. Archivists on both sides of the debate have strong feelings about this. George Blood supports the JPEG 2000 standard in large part because of the official documentation behind it, as does Jim Lindner. Linder doubts the validity of both FFV1 and the actions FFV1 advocates are taking to standardize it:

> FFV1 is not a standard. It is so not a standard that they have tried to do an end run around the normal organizations that create these standards to establish their own standard. I think that says [allot]. The people who work at MPEG/SMPTE/ISO are really smart, have done it for many years and know what they are doing… but instead the plan is to throw it against the wall and for a group that does not really set these types of standards to set one? Why? Well I'm sorry I have no respect for that. Kind of like I will take my marbles and go home if you don't play the game I want to play. Really an immature attitude and not the kind of 'standard' that any organization that uses standards should accept.[23]

Lindner's rhetoric here shows the trust he puts in standards organizations. JPEG2000 supporters value a certain "by the book" approach to choosing their codec, and Lindner finds the act of the CELLAR working group to write their own FFV1 and Matroska standardization to be an act of immaturity. He shows distrust in the Internet Engineering Task Force as a standards body, as their work has dealt primarily with standards for internet protocols, rather than file formats or preservation activities. Since the IETF has, up until now, not been the standards organization of

[23] Lindner

choice for file format specifications and preservation protocols, Lindner appears to doubt the validity of FFV1 as a standard.

Meanwhile, FFV1 supporters are wary of JPEG 2000 for how its documentation is able to be accessed. According to Dave Rice, "JPEG 2000 is technically an open format but there's a cost and constraints to access the documentation. FFV1 is openly documented in an open standards org, so the documentation and entire process and conversation behind it is open to research."[24] Though JPEG 2000 is standardized, the standard is locked behind a paywall. If one wishes to view it, they must have to either be a paying SMPTE subscriber, or pay an exorbitant fee to access each one of the eight parts of the standard. Additionally, the standardization process is not transparent. If a standard has been revised, like JPEG 2000 has, only the most recent version of the standard's documentation is available at all, as SMPTE removes all previous versions from their website. This can be a hindrance to researching the history of JPEG 2000 as it is impossible to see what changes have been made to the format or any documentation about why the changes were made.

When considering what standards body to body to work with to legitimize FFV1 and Matroska, members of the CELLAR working group initially considered ISO and SMPTE, but decided on IETF, the Internet Engineering Task Force. They found that, "…payment for access to standards runs counter to the ideals of the […] project. Within the context of digital preservation it is imperative for file formats to be well-disclosed, understood, and controlled within an archival setting; specification paywalls provide an obstacle to this objective."[25] As

[24] Rice, "Re: Thoughts on video codecs?"
[25] Ashley Blewer, Tessa Fallon, and Dave Rice, "Conch – Appendix on Standardization Exercises," 02 Mar. 2015.

FFV1 is primarily used in archival workflows that utilize open source tools, it follows that the archivists working to standardize the format want it to be as well-documented and easily accessible as possible. Anyone can join the CELLAR mailing list and be a part of the discussion around standardization and adding new features to FFV1. Even though it is not formally a standard yet, extensive documentation of FFV1's specifications is freely available, including through FFmpeg's website.

The ideologies of the two factions can be distilled down to this issue: FFV1 supporters are also advocates for free and open source tools and creating a community of archivists and developers that work together to make the resources they use for preservation better. Meanwhile, JPEG 2000 supporters like the perceived security of an official standard and proprietary software and hardware. These attitudes come into play at the institution-level too, as what the institution values may lead to them choosing one codec over the other.

Institutional Considerations

Ultimately, when an institution is deciding which file format to implement for their preservation masters, there are a multitude of facets to consider. According to George Blood, "Independent of context, the debate is meaningless. Neither will win on its own merits… There are much bigger institutional issues at play."[26] Both JPEG 2000-advocate George Blood and FFV1-proponent Peter Bubestinger suggest a "try it before you buy it" approach, seeing first hand which format works well within the institution's pre-existing infrastructure and workflows.[27] Some factors to consider are whether or not the institution uses predominantly open

<http://www.digitalmeetsculture.net/wp-content/uploads/2015/04/MediaAreaConch_Appendix_Standardization.pdf>

[26] Blood

[27] Bubestinger, "Re: Thoughts on video codecs?"

source or proprietary software, and how comfortable employees are using the command line as opposed to graphical user interfaces.

Time and money are also two factors involved with this decision, as FFV1 is cheaper to use and faster at encoding than JPEG 2000. In addition to testing both codecs himself, Bubestinger says the Austria Mediathek also, "…chose FFV1 for pragmatic reasons, after speaking to institutions that were using JPEG 2000-lossless (mostly in MXF): We simply couldn't afford wasting so much time and money on problems that are just not present in FFV1."[28] Meanwhile, at an institution like the IFI Irish Film Archive, which is just starting to implement a high-volume digitization project of their holdings, Kieran O'Leary chose FFV1 for their workflows, saying, "We are a small archive with a limited budget and a small amount of staff," and he cites the open documentation of FFV1 and its ability to store slice-level checksums as other factors in this decision.[29] O'Leary is also advocating for FFV1 for their film scans, as well as for their digitized analog videos. Additionally, some prefer FFV1 for its speed and easy implementation within FFmpeg. According to Dave Rice, the archivist at City University of New York Television, "Selecting it at CUNY was mostly because we use FFmpeg for nearly all media processing and we can work with FFV1 with FFmpeg much faster to encode, decode, and validate."[30] As FFmpeg was already a large part of CUNY's workflow, it made sense for them as an institution to use FFV1 as their preservation format.

To conclude, the choice of using JPEG 2000 or FFV1 is not just a matter of efficiency, quality, or cost, but the core values of the archival institution. As Rice notes, "JPEG 2000 use

---

[28] Ibid.
[29] Kieran O'Leary, "Introduction to FFV1 and Matroska for Film Scans," 07 Oct. 2016. <https://kieranjol.wordpress.com/>
[30] Rice, "Re: Thoughts on video codecs?"

and support is mostly associated with proprietary systems and software whereas FFV1 use and support is mostly from open source communities, so the decision on whether to use FFV1 or JPEG 2000 should consider what other users, systems, and communities that you'd like to work with."[31] As exemplified in the disagreement over standardization organizations, JPEG 2000 communities and FFV1 communities represent different ideologies. Though both groups are concerned with choosing a lossless format that is an accurate representation of the original video content, they differ in terms of their respective values. JPEG 2000 users prefer the security of using a format that is endorsed by trusted international standards bodies and supported by large hardware and software-producing companies. Meanwhile, FFV1 advocates value accessibility and the transparency in the process of creating the standard more than the standards organization itself. To this end, if your institution choses one codec over the other, they are choosing a specific user community to be involved with, and the inherent ideology underlying that community.

[31] Ibid.

Works Cited

Blewer, Ashley and Dave Rice. "Status of CELLAR: Update from an IETF Working Group for Matroska and FFV1," *Mediaarea.net*, 03 Oct. 2016. Web. 01 Dec. 2016. <https://mediaarea.net/Events/PDF/2016-10-03_iPRES_Status_of_CELLAR.pdf>

Blewer, Ashley, et al., "Conch – Appendix on Standardization Exercises," 02 Mar. 2015. <http://www.digitalmeetsculture.net/wp-content/uploads/2015/04/MediaAreaConch_Appendix_Standardization.pdf>

Blood, George. Personal interview with Savannah Campbell. 28 Nov. 2016.

Bubestinger, Peter. "Re: Thoughts on video codecs?" Message to Savannah Campbell, 28 Nov. 2016. Email.

Bubestinger, Peter, et al. "Comparing Video Codecs and Containers for Archives," 13 Aug. 2015. Web. 18 Nov. 2016 <http://download.daswerkstatt.com/pb/mthk/info/video/comparison_video_codecs_containers.html>.

Elzinga, Jason and Keith Feenstra. "JPEG 2000: The Next Compression Standard Using Wavelet Technology," 04 Dec. 2001. Web. 01 Dec. 2016. <http://faculty.gvsu.edu/aboufade/web/wavelets/student_work/EF/how-works.html>

Fleischhauer, Carl. "Comparing Formats for Digitization," *loc.gov*. 03 Dec. 2014. Web. 29 Nov. 2016. <http://blogs.loc.gov/thesignal/2014/12/comparing-formats-for-video-digitization/>

Lewetz, Herman. "The MediaConch Presentation at IASA," *mediaarea.net*. 22 Oct. 2015. Web. 3 Dec. 2016. < https://mediaarea.net/MediaConch/2015/10/22/iasa-presentation/>

Lindner, Jim. “Re: Thoughts on video codecs?” Message to Savannah Campbell, 26 Nov. 2016. Email.

Nidermayer, Michael. “FFV1 Video Codec Specification,” ffmpeg.org. Web. 20 Nov. 2016. <http://www.ffmpeg.org/~michael/ffv1.html>

O’Leary, Kieran. “Introduction to FFV1 and Matroska for Film Scans,” 07 Oct. 2016. Web. 28 Nov. 2016. <https://kieranjol.wordpress.com/>

O’Leary, Kieran. “Re: Thoughts on video codecs?” Message to Savannah Campbell, 28 Nov. 2016, Email.

“Overview of JPEG 2000,” *jpeg.org*. Web. 20 Nov. 2016. <https://jpeg.org/jpeg2000/>

Dave, Rice. “Re: FFV1 vs other formats for preservation,” *Archivematica Google Group,* 24 Sept. 2012. Web. 01 Dec. 2016. <https://groups.google.com/forum/#!topic/archivematica/HulV96gJ0go>

Rice, Dave. “Re: Thoughts on video codecs?” Message to Savannah Campbell, 27 Nov. 2016. Email.

Snyder, James. “JPEG 2000 in Moving Image Archiving,” 13 May 2011. Web. 20 Nov. 2016. <http://www.digitizationguidelines.gov/still-image/documents/Snyder.pdf>

“Sustainability of Digital Formats Planning for the Library of Congress: FF Video Codec 1,” 23 Dec. 2015. Web. 01 Dec. 2016. <http://www.digitalpreservation.gov/formats/fdd/fdd000341.shtml>