Federal Agencies Audio Visual Digitization Working Group

Interstitial Error Study

Volume I. The Study Report

Prepared
By

AUDIOVISUAL PRESERVATION SOLUTIONS

350 7th Ave.
Suite 1605
New York, New York
10001
Tel: 917.475.9630 | Fax: 866.264.4275

www.avpreserve.com

Submitted April 8, 2013

TABLE OF CONTENTS for *Volume I. The Study Report*

LIST OF CONTENTS for *Volume II. Appendixes*

Background

This two-volume report is one product of a project commissioned by the Federal Agencies Digitization Guidelines Audio-Visual Working Group (see www.digitizationguidelines.gov). The aim of this project is to advance user capabilities to test and measure audio digitization systems that are deployed in archival or preservation settings, and the effort has also produced the guideline *Analog-to-Digital Converter Performance Specification and Testing.*[1] The author of this document and principal investigator for this project is Chris Lacinak of AudioVisual Preservation Solutions, a special consultant to the Working Group.

This first volume is largely devoted to a report on the field test carried out in the summer of 2012, with final editing taking place in early 2013. The second volume, titled *Appendixes*, provides background information as well as re-presenting sections from a 2011 report. The *Appendixes* were completed in September 2012.

One area of investigation is the development of a solution for detection and reporting of 'interstitial errors'. The term ‘interstitial error’ was coined by this writer in a white paper published in January 2010, titled *Digital Audio Interstitial Errors: Raising Awareness and Developing New Methodologies for Detection*[2]. Interstitial errors consist of lost or altered samples within the recorded file, resulting in the loss of content and integrity. These errors -- often very momentary -- result from a failure in the chain of digital data, i.e., in the handoff from the analog-to-digital converter (ADC) to the digital audio workstation (DAW), and in the DAW's writing of the file to a storage medium. In February 2010, the Working Group held a meeting which addressed the topic of audio system evaluation. During this meeting anecdotal support was provided by Working Group members, identifying interstitial errors as a problem in need of address. As a result of this meeting, three areas were identified for further research and development, one of which was interstitial errors.

In March 2011, an interim report was published through the Working Group[3] addressing the areas of interest with regard to audio system evaluation and providing an outline of next steps. For interstitial errors, the report proposed the following next steps: a) more clearly specify the requirements and features of a tool for detecting and reporting errors; b) establish a workplan and estimated level of effort for developing such a tool, and c) perform a wide-reaching survey to assess the level of awareness of the issue and interest in a tool for detection and reporting.

In 2011 and 2012, these next steps were taken. The survey results yielded a total of 83 respondents, 56 of which either had first hand experience or had encountered interstitial errors. 75 respondents expressed interest in having a tool to address the issue. However, there were several voices that expressed opinions to the contrary. Their comments ranged from statements that held that interstitial errors were a non-existent issue to statements that

---

[1] http://www.digitizationguidelines.gov/guidelines/digitize-audioperf.html
[2] http://www.avpreserve.com/wp-content/uploads/2010/01/Digital_Audio_Interstitial_Errors.pdf
[3] http://www.digitizationguidelines.gov/audio-visual/documents/FADGI_Audio_EvalPerf_Report.pdf

associated these errors with inadequate resources (e.g., lack of operator expertise, need for top-of-the line DAWs).

Based on the 2011 report and the survey that followed, the Working Group decided to launch an exploratory field study to begin the process of quantifying the issue and putting it into perspective. The field study was performed during the spring and summer of 2012, and this report details the study's process and findings.

Introduction

The field study described in this report collected 171 samples representing outputs from three devices, each of which was operational for three to four weeks. The systems tested were located in federal government agencies that regularly carry out audio digitization projects and operate audio preservation and digitization facilities.

The study was intended to address the questions, "Does this problem exist, do these errors occur?" more than "What are the causes and cures?", although the latter question was a subtext throughout the investigation. And as noted above, limited resources constrained the extent of the field study: we gathered a relatively small number of sample files from three workstations, inhibiting us from having a basis for a statistically meaningful finding. Nevertheless, we wished to (a) assess the extent to which interstitial errors exist and the frequency and severity with which they occur, and (b) model a prototype test setup and method for detecting and reporting errors. In addition, as reported in the section titled *Unanticipated Findings*, the prototype test setup incidentally (and helpfully) revealed other operational issues in the respective recording facilities.

Given the intermittent nature of interstitial errors, this extent of testing offers a useful and instructive glimpse of the issue, and the Working Group feels that this study provides the preservation community with a good starting point for future investigations. Readers should be cautious, however, when extrapolating from this limited sample. In the following sections, given the study's limited and potentially unrepresentative sample set, we have refrained from associating the specific results with the participating agencies and we have not identified the DAWs that were used.

The setup and method borrowed a concept proposed in the Working Group report from March, 2011. The concept can be seen in Figure 1. Once AVPS compiled and tested the prototype, it was deployed in three active digitization facilities. The prototype system was developed to be integrated into working systems within test sites, and as such was designed in such a way that it would not alter the integrity of the test site's signal path or workflow used for recording to their digital audio workstation (DAW).

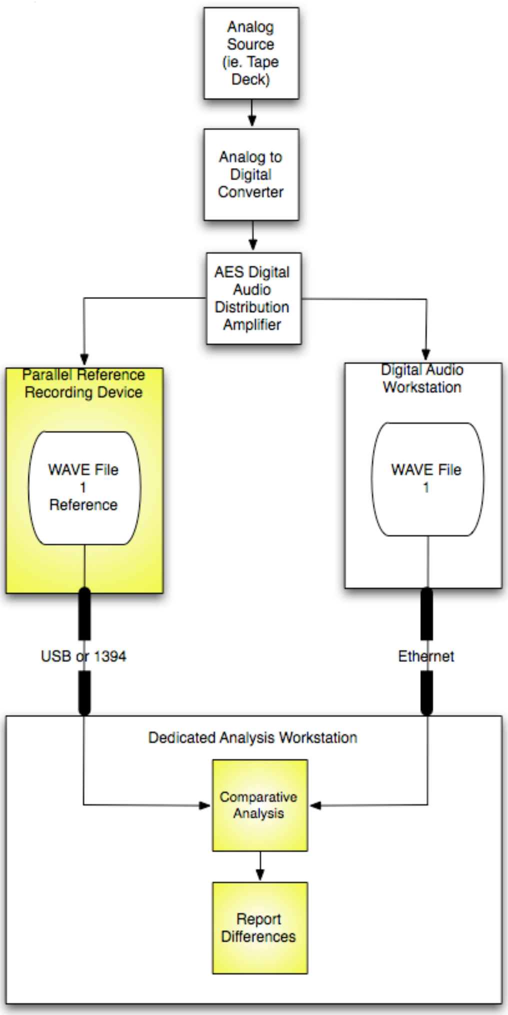

Figure 1: Concept Used for Prototype System

The prototype system was used to create an exact duplicate of the AES digital audio output from the test site's analog-to-digital converter (ADC). One digital audio stream was routed to the digital audio workstation according to typical protocols. The other digital audio stream was routed to a standalone file-based direct-to-disk recording device (DDR).

Project equipment was installed and tested by AVPS at each site, and the local technicians at the agencies collected data throughout each test period. Data was collected over a five-month period, with three to four week test periods in each location. At the end of each data collection period all file pairs and associated documentation were sent to AVPS. Once at AVPS the files were analyzed. The result is a set of data speaking to the extent to which interstitial errors may be a problem and the frequency with which they may occur.

Prototype Test Setup

The prototype test setup consisted of the following equipment:

RME Fireface UFX digital audio interface
The RME Fireface UFX was selected due to its ability to serve as both digital distribution amplifier and file-based recording device. The unit accepts an AES digital audio input, and creates a duplicate digital audio stream. The UFX passes one digital audio stream through its AES digital audio output and routes the other stream to a front-mounted USB port purpose-built for direct recording to a hard drive or flash drive. The unit requires only its own firmware to write audio files to the connected USB drive.

2 x 2TB Olixir Data Vault external hard drives
One of these drives was used for collecting files from the DAW after they were recorded to the DAW as usual. The other drive was used to connect to the USB port on the front of the UFX for direct-to-disk recording.

AES cable
An additional AES cable was provided to make the link between the UFX and the host institution DAW.

At the beginning of each test period, AVPS installed and tested the prototype test setup at each site. The technician was trained in use of the test equipment, and a step-by-step procedure manual was provided to offer additional instruction and support[4].

The equipment was installed such that all controls were easily accessible and that standard operating procedures were able to be followed with minimal disruption. The technician was also asked to provide basic documentation via a provided template[5].

We encountered a few anomalies with the test setup as implemented in these three sites. This topic is discussed in Appendix D (all appendixes are in Volume II).

Process and Methodology

After turning on and initializing all equipment, technicians were asked to prepare a recording for digitization as normal. The technician would then start recording on both the DDR and DAW at approximately the same time, and then start the playback equipment. At the end of a recording, technicians were asked to stop both the DAW and DDR at approximately the same time. This process was repeated for as many recordings as were made during a work session.

---

[4] See Appendix A (all appendixes are in Volume II)
[5] See Appendix B

Each institution observed its normal file-naming protocol for recordings created on the host DAW. The DDR uses a fixed file naming protocol that does not allow user modification. Technicians were asked to record the file names of their DAW recordings and the corresponding DDR files on the provided log sheet.

At the end of each day's work, the technicians were asked to copy the files created on the DAW to a dated folder on the provided DAW external hard drive, and to move the day's DDR files to a dated file on the DDR hard drive. This daily procedure was repeated throughout the test period.

The test sites created recordings in mono, dual mono, and stereo track configurations. All collected files were 24bit/96kHz WAV files. Each location utilized a unique audio hardware and software chain as well as unique file naming practices.

Given the nature and scope of this study, specifications and configuration settings for each system were not documented as part of the study. At a high level, the DAWs represented include:

- Turnkey and non-turnkey
- Lower cost and higher cost
- Multiple operating systems

Timeline for Each Test Site

The test was carried out in periods lasting from 16 to 20 working days at each test site.
In the case of the first test site, a clocking error was discovered in the DAW files upon analysis rendering a large majority of the files from that test period invalid for the purpose of this study, and a second collection period was performed.

Test Site	Begin Date	End Date	# of usable file pairs collected	Recording hours (hh:mm:ss)
Test Site 1A	03/23/2012	04/19/2012	10	07:20:56
Test Site 2	04/26/2012	05/17/2012	56	27:02:42
Test Site 3	05/30/2012	06/21/2012	33	23:22:44
Test Site 1B	07/02/2012	07/26/2012	72	53:02:17
TOTAL	03/23/2012	07/26/2012	171	110:48:39

Figure 2. Testing dates.

Analysis Methodology

At the end of each test period all files were sent to AVPS. Upon receipt, MD5 checksums were created for all files. The data was then transferred to the AVPS DAW and checksums were verified.

File pairs were identified according to the logs submitted by each test site. The DAW and DDR files for each pair were opened in WaveLab and brought into a montage. The files were synchronized with sample level accuracy to a common point near the beginning of the files. With the files synchronized at the start, the DDR file from each pair was polarity inverted, so that a summing of the two files would, in the absence of any interstitial errors, create a null outcome, or digital silence. Each montage was rendered to a new 24bit/96kHz WAV file, referred to as the *delta* file. The delta file was normalized to 0dBFS to make any variances visually apparent, and then visually inspected.[6]

If no variation from null was seen on the waveform, it was rated a "pass," with no interstitial errors. If signal was identified, it was flagged and the sample count of the error was documented for further inspection. Further inspection consisted of going back to the montage and reviewing the files at the documented sample count. Upon confirmation of an error on the DAW file, the file was marked as a "fail". The number of missing samples was also recorded, both at the site of the error and again at the end of the montage, to determine whether multiple errors had occurred. If there was a discrepancy between the number of missing samples at the end of the file and the number of missing samples in the single error, further analysis was performed to identify the additional errors. This further analysis involved re-synchronizing the audio waveform after the point of the error and continuing to perform the null test. This realignment of the waveform can be seen in Figure 6.

Findings

Unanticipated Findings

The installation of the test setup unexpectedly revealed issues that were unrelated to interstitial errors. After each installation and training, the whole system was tested to make sure that everything was functioning properly. To do this, content was recorded according to protocol in the procedure manual and a null test was performed immediately. Generating a delta file with all digital silence was sufficient evidence of a successful installation. If the delta file did not contain digital silence it was proof that something was not functioning correctly. In

---

[6] Normalization is the process of increasing the maximum peak value in an audio file to a particular audio level. In this case it was 0dBFS, the maximum value on a digital peak meter. In this scenario, if there were any value that was greater than null it would be made to equal 0dBFS and therefore it would become visually apparent.

three out of four installations,[6] when this initial test was performed, it resulted in a delta file indicating that there were issues calling for further investigation.

In the first install at Test Site 1, the principal investigator found that the DAW clock was set to "internal" instead of slaving to the AES input. Once he and the system operator changed this setting, the testing produced correct delta files. The operator reported that this was a new DAW setup that had not been used prior to this study, which meant that this operational error had not caused any defects in files produced prior to the field study. However, after three days of test recording this setting automatically and inadvertently switched itself back to "internal." In this mode, the DDR clock slaves to the ADC while the DAW clock does not, a circumstance that ruled out a null test on the pairs of files that were produced. Thus a second visit was scheduled to capture additional files. What is notable is that this issue would have gone unnoticed without performing this analysis, showing evidence of the additional value of this simple test. During the second round of testing the DAW clock setting was monitored closely and this issue did not occur.

During the install at Test Site 2 problematic delta files were also found. Further investigation revealed that the DAW clock was behaving erratically. After extensive troubleshooting and assistance from the DAW manufacturer's tech support, it was found that the firmware and software was years past due for an update. The tech support engineer disclosed that earlier versions (including the one being used prior to updating) were extremely problematic. After updating of all firmware and software successful delta files were achieved. Again, although this test was not meant to identify these types of issues, until this comparative analysis test was performed the problem was unknown.

During the second install process at Test Site 1 the delta files produced showed evidence of interstitial errors. Through several tests, this occurred in each file, and several times within a file. Through routine troubleshooting, the DAW was restarted. Upon reboot, the principal investigator observed an antivirus application loading. Soon, the team determined that this software had been installed by the organization's IT staff in between the end of the first install and this install. Once booted, exploration of the antivirus application settings revealed a preference for "scan files while writing to disk." Disabling of this feature resulted in correct delta files. Ultimately the team disabled and shut down the antivirus application completely, and this event was used to inform policies regarding the use of anti-virus applications and related issues in the audio preservation lab. In this case, the organization had been using their DAWs with this installed for some time and it was not until we ran this test that it was identified.

---

[6] Test Site 1 went through two separate rounds of testing.

Interstitial Error Findings

Test Site 1 (A&B)

Over 17 days, 82 valid pairs of files were collected, representing just over 63 hours of recording time. One interstitial error was detected. This file did not contain missing samples, but rather a rearrangement in the first 763 samples. The sample count was correct but the ordering of the first 763 samples was not correct.

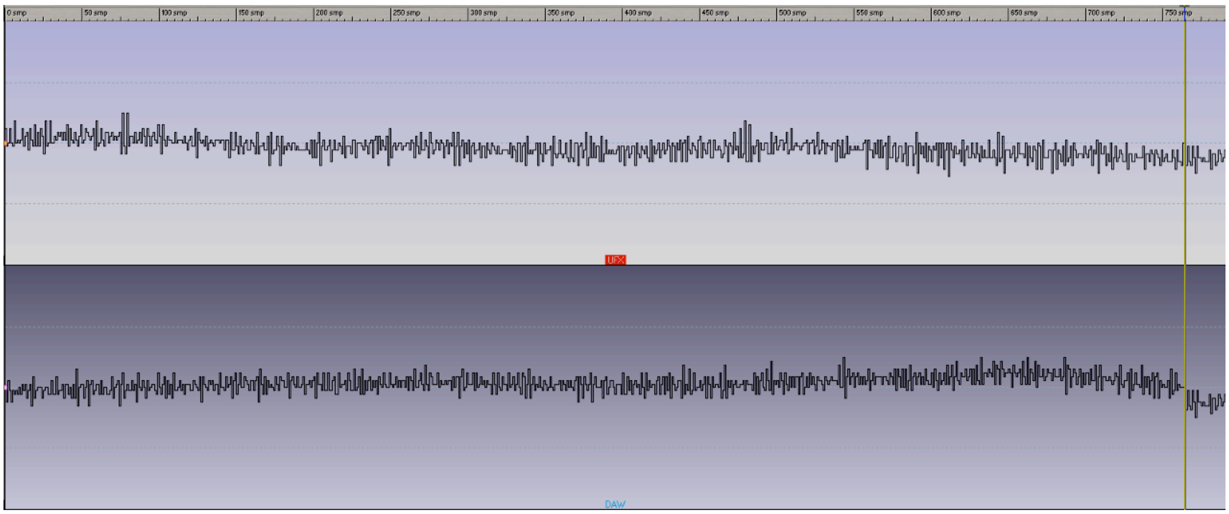

Figure 3: Demonstrating the interstitial error in which the correct number of samples were written to disk, but the first 763 samples are not the correct data. The vertical line represents the point at which the correct writing of samples begins.

This results in a failure rate of 1.2%. At 96000 samples per second this equates to approximately .008 seconds of time in approximately 23 hours and 20 minutes of recording time, or 83,988 seconds. This represents a .002% loss of content, or approximately 1 hour for every 46,700 hours of recorded content (consisting of approximately 14,151 errors based on the average of 1 error every 3.3 hours).

Test Site 2

Over 15 days, 56 valid pairs of files were collected, representing just over 27 hours of recording time. No interstitial errors were detected.

Test Site 3

Over 13 days, 33 valid pairs of files were collected, representing over 23 hours of recording time. 6 DAW files were found to have interstitial errors, totaling 7 errors (one file had two interstitial errors). This results in a failure rate of just over 21%, or an average of 1 error every 3.3 hours of recording. The number of samples dropped in a given interstitial error ranged from 4,100 to 53,200, with an average of 24,750 samples lost. There were a total of 148,500 samples lost. At 96,000 samples per second this equates to approximately 1.5 seconds of time in approximately 23 hours and 20 minutes of recording time, or 83,988 seconds. This represents a .002% loss of content, or approximately 1 hour for every 46,700 hours of recorded content (consisting of approximately 14,151 errors based on the average of 1 error every 3.3 hours).

Examples of Errors

Simple Comparison

Figures 4, 5 and 6 below provide a side-by-side comparison of a DAW file containing an interstitial error and the DDR file that does not have an error. They are presented as waveforms in a timeline. For clarity in these illustrations, the polarity of the DDR file has not been reversed, as it is when the null test is performed. You will notice that the waveforms are the same from the beginning up until a certain point after which they differ. The point at which they differ is where the interstitial error occurs; as will be shown in the next section, this error point usually indicates that a series of samples are missing from the DAW file.

Figure 7 illustrates the alignment at the beginning of a file in contrast to the end where a shift forward can be seen in the DAW file caused by the loss of samples.

Figure 4: DAW file with interstitial error compared to the intact DDR file

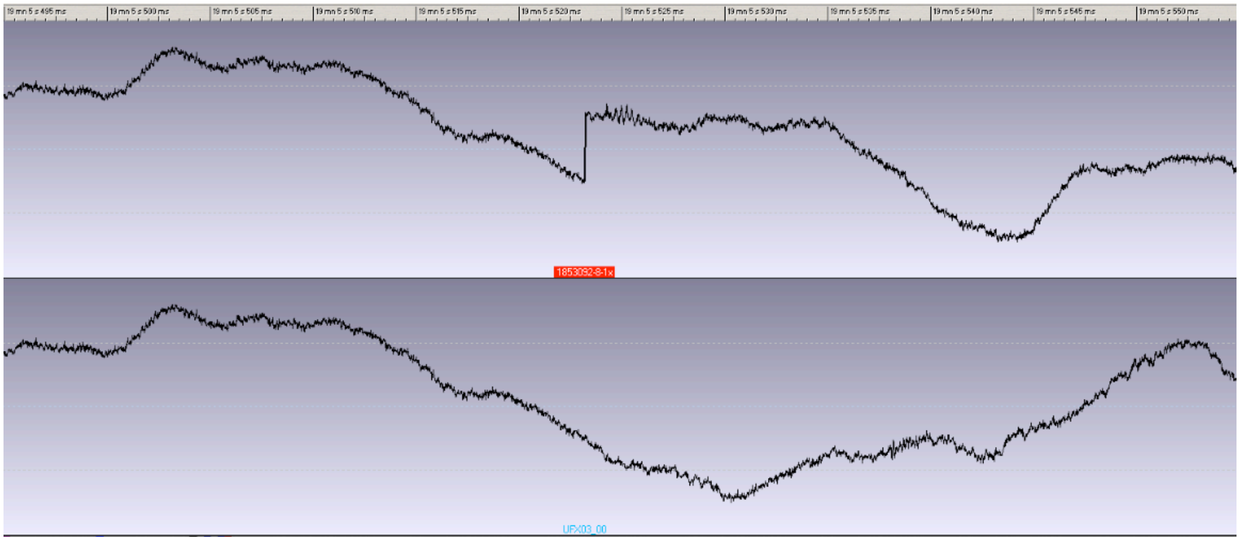

Figure 5: DAW file with interstitial error compared to the intact DDR file

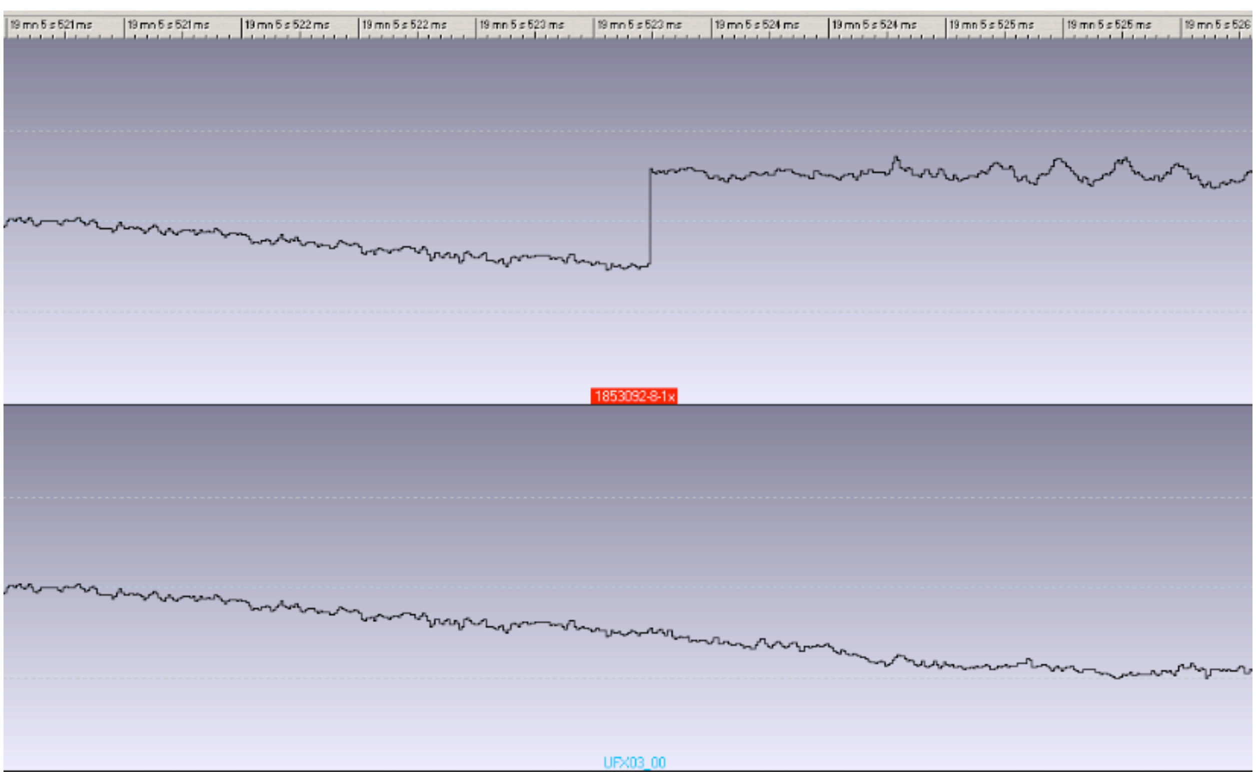

Figure 6: DAW file with interstitial error compared to the intact DDR file

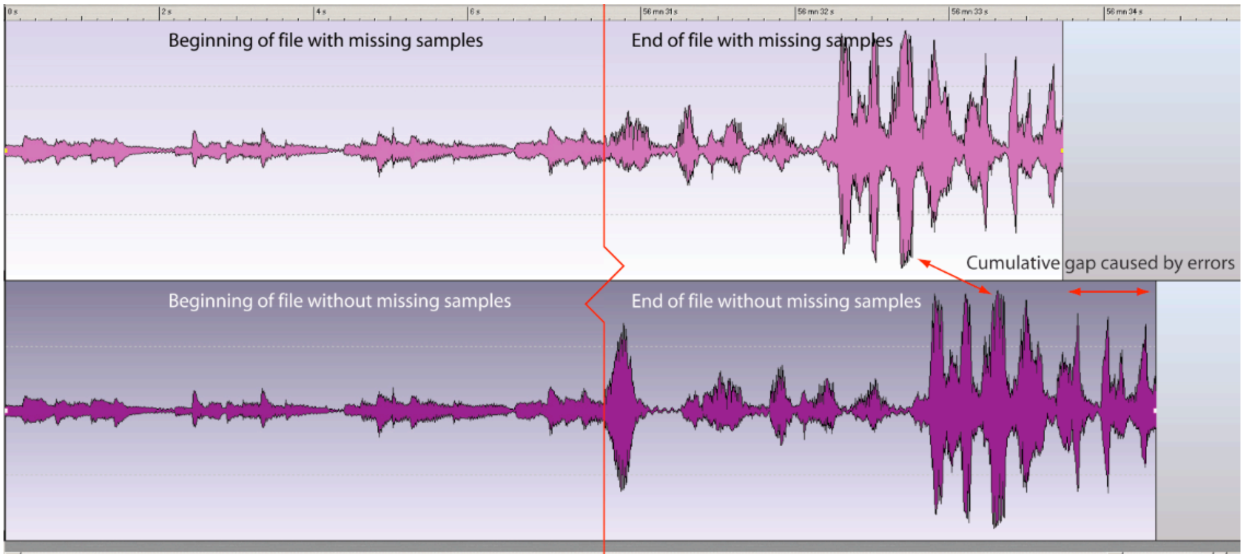

Figure 7: Demonstrating the beginning and end of the DAW and DDR files. Note that they are synchronized in the beginning and that by the end they are out of sync due to missing samples in the DAW file.

Comparison Including the Null Test

Figures 8 and 9 illustrate how the waveforms were analyzed in this field test. The two waveforms at the top of each figure are similar to what is shown in figures 4, 5, and 6, except in these figures the polarity of the second (DDR) waveform has been reversed. The third waveform in these sets shows the DAW file but aligns all of its content with the DDR file. Thus, the gap represents the missing samples. The last waveform is the delta file, comparing the DDR file with the DAW file; the "null" results are the flat line while the samples missing from the DAW file are represented by a waveform contour.

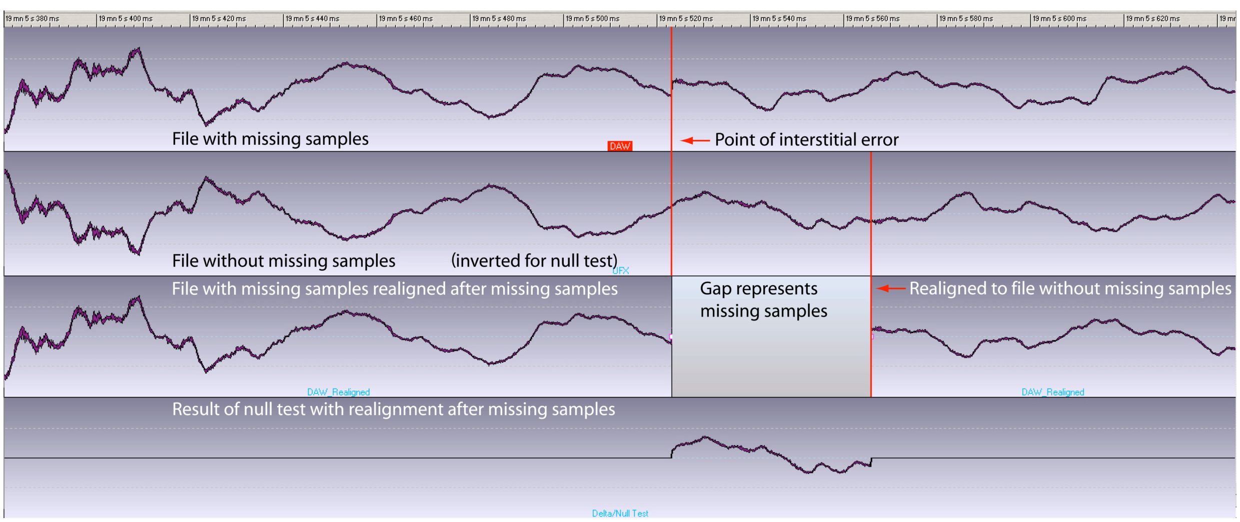

Figure 8: Demonstrating the original DAW file, the DDR file and the Delta file along with a DAW file showing a gap where the missing samples are.

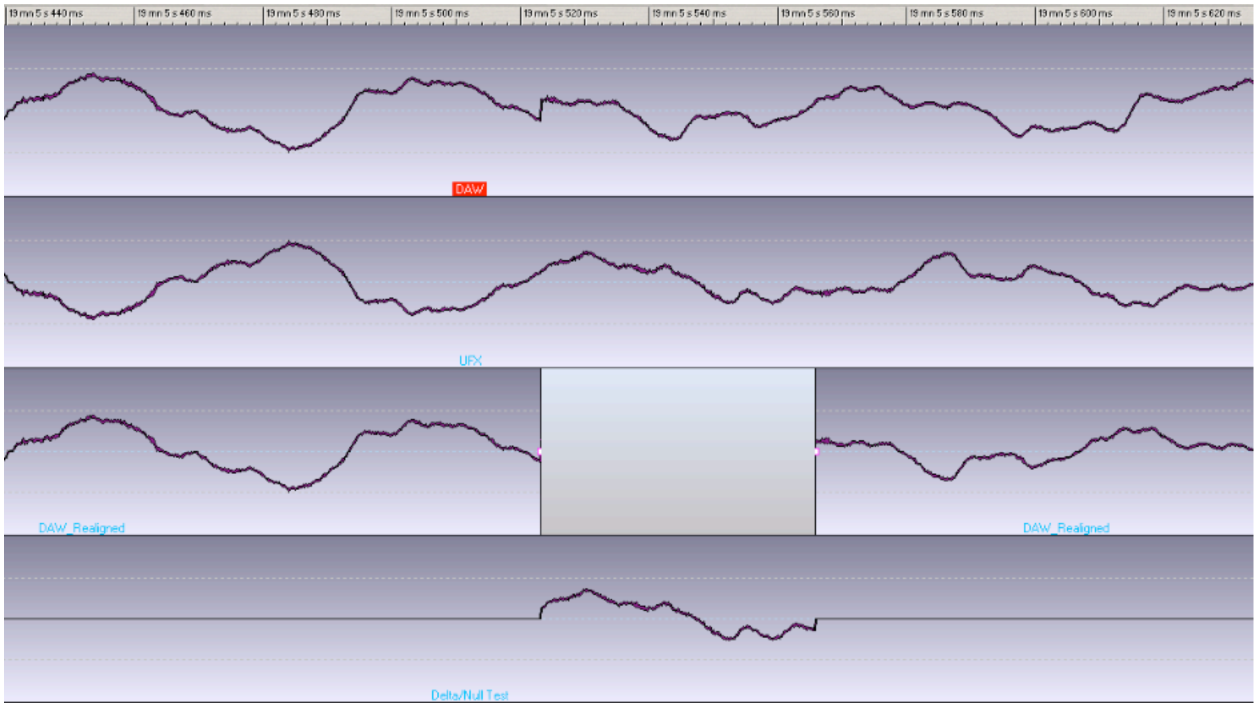

Figure 9: Demonstrating the original DAW file, the DDR file and the Delta file along with a DAW file showing a gap where the missing samples are.

Sound Excerpts of Files

The URLs that follow link to excerpts from pairs of audio files that exhibit interstitial errors documented during this study. Each pair contains one file with an error and another file without the error. All of these examples represent missing-sample errors. We have not included an example of the files in which the error consisted of rearranged samples. This error occurred in a section of a recording that contains only very low level audio and nothing can be heard.

201206081082 File Set
http://www.avpreserve.com/interstitialerrorsamples/201206081082_DAW.wav
http://www.avpreserve.com/interstitialerrorsamples/201206081082_DDR.wav

201206121086 File Set
http://www.avpreserve.com/interstitialerrorsamples/201206121086_DAW.wav
http://www.avpreserve.com/interstitialerrorsamples/201206121086_DDR.wav

Overall Analysis

Overview and Caveats

As noted in the introduction, we do not feel that our 171 files from three digital audio workstations provide enough data to support statistically valid findings. Ideally, one would want a larger number of files. In addition, the three DAW units were from different manufacturers; a more thorough study may have compared multiple instances of DAWs of the same make and model. We were working with limited resources and also considered the "imposition" factor. We are grateful to the three organizations that permitted us to use their facilities and staff time to carry out the field test; to do more at this time would have imposed on those facilities and staff to a greater degree.

The preceding paragraph is intended to offer a caveat for the summary that follows. We have calculated some percentages and totals and extrapolated against a large number of hours that a lab might produce. **We are confident that the field tests demonstrate that interstitial errors occur in the normal operation of a recording facility. We are less confident of the suggested level of error across a large production run. To be compelling, such an extrapolation would require a statistically meaningful basis.** It is also worth noting--as reported in the section titled *Follow-on Testing* below--that one of the three participating organizations reported having run a series of follow-up field tests, producing 40 recordings, none of which yielded interstitial errors. (These results are not included in our summation.)

Analyzing all test sites, based on file count the failure rate is 4.7%, or 1 error every 13.85 hours of recording. Based on time, the loss factor is .0005%, or 1 hour lost for every 200,000 hours of recording (consisting of approximately 14,440 errors based on the average of 1 error per 13.85 hours of recording). It should also be noted, as mentioned in *Unanticipated Findings,* that during installation of the system in the second round of testing at Test Site 1 multiple interstitial errors were detected in every file produced by the DAW. The source of the issue was discovered as being anti-virus software and remedied. Had the errors not been caught upon installation, it is likely that many more files would have had interstitial errors, greatly altering the results.

Figures 10 and 11 summarize the results from this study. The full details can be seen in Appendix C (all appendixes are in Volume II).

Test Site	Usable File Pairs Collected	DAW Files w/Errors	Total Time (hhh:mm:ss)	DAW Failure Rate
Test Site 1	82	1	60:23:13	1.2%
Test Site 2	56	0	27:02:42	0%
Test Site 3	33	7	23:22:44	21%
Total	171	8	110:48:39	4.7%

Figure 10. Test result analysis by file

Test Site	Total Time (seconds)	Seconds Lost	DAW Failure Rate
Test Site 1	217393	.008	.000004%
Test Site 2	97362	0	0%
Test Site 3	84164	1.5	.002%
Total	398919	1.508	.0005%

Figure 11. Test result analysis by time

Error Location in Files

During the course of the investigation the question of whether errors occur more frequently at the beginning, end, or middle of files arose. An investigation into this yielded findings that demonstrated no clear pattern. Following this we looked into whether the errors consistently occurred relative to the length of the files. The graph in figure 12 plots the location of errors in terms of the duration of a file, expressed here as a percentage of the file's length. The graph indicates that, at least in this limited test, the errors tend to be distributed throughout the files, not favoring any particular location.

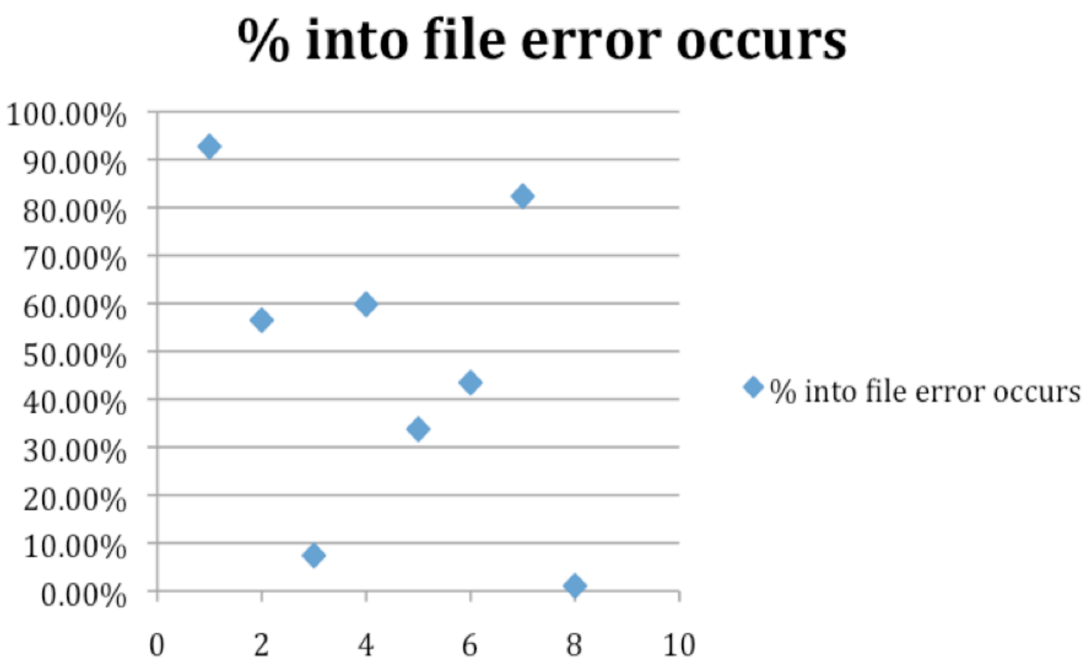

Figure 4: Location of errors in files, expressed as percentage of the duration of the recording

Additional Errors

This is a summary of information provided in appendix D. When designing this field test, we had expected the files produced in the direct-to-disk recording (DDR) system, and its associated storage devices, to be free of interstitial errors. However, we found three DDR-file errors, two from systems with USB flash drives and one with external hard drive, and do not have an explanation.

Meanwhile, the workstation in one of the test sites consistently produced an error that is not exactly *interstitial*, as that term is used in this report. The DDR files from this site lost the "values" for 512 samples (less than 1/100 of a second) at the very beginning and end of each recording. This may have resulted from a unique feature of that site's setup. At two of the sites, we employed the preferred approach, in which the UFX passes one digital audio stream through its AES digital audio output to the DAW while simultaneously routing the other stream to a front-mounted USB port for direct recording. At the problem site, the hardware configuration precluded this setup and, instead, the split occurred within the IO for the DAW. It is possible that the 512-sample drops resulted from the signal management of the DAW.

Follow-on testing

Following the field test described in this document, one of the three organizations offered to perform some additional field testing on their own, looking at files from some additional digital audio workstations using the same technology employed in the main test. During December 2012 and January 2013, the organization produced 30 files from one digital audio workstation and about 10 on another. They reported that they did not identify any errors in any of these files. They hope to do some additional testing in the coming months.

Conclusions

The purpose of this study was to look at the extent to which interstitial errors exist and the frequency with which they occur. The limited testing performed revealed that these errors do in fact exist and it provides some very preliminary findings on the frequency with which they may occur.

An unanticipated finding of this study was the potential value of using a test system like the one employed here to detect problems other than interstitial errors. As reported, we identified issues including clock performance and equipment malfunction in two of the three test sites. Also, a badly behaving anti-virus application was immediately found to be producing interstitial errors and disabled before continued testing. This issue had not been identified prior to our testing.

Based on the cumulative findings, we believe that interstitial errors are an issue worthy of serious consideration in an archival context, where matters of authenticity, integrity and quality are concerned. They are, in fact, a problem to be addressed in the archival community as well as by the Working Group and we have framed the following section on next steps for community consideration.

Next Steps

Ideas for next steps have emerged from the March 2011 report to the Working Group,[7] the field test, and a variety of informal conversations. What follows is a list of possible activities together with summary comments representing talking points from these conversations and thoughts on moving forward.

- **Additional Investigation of the Problem: Compile a Statistically Valid Study**

  This would consist of a study with a larger number of DAWs, including multiple instances of the same make and model, and the production of a larger number of files.

  We feel that this would be an effort with limited value. Our field study and other prior evidence, as summarized in the March 2011 report, convincingly indicates that interstitial errors occur, albeit at random intervals.

---

[7] http://www.digitizationguidelines.gov/audio-visual/documents/FADGI_Audio_EvalPerf_Report.pdf

- **Cause Analysis: Determine the Operational Elements Encountered in DAWs that Cause Interstitial Errors**

  This would consist of a set of experiments to further investigate causes. One approach would be to give each postulated cause its own set of experiments. Although the execution of a full range of well-designed experiments would be resource intensive it may be that some very interesting things could be discovered with some well-selected experiments of limited scope and depth.

  This writer believes that interstitial errors result from multiple and variable causes which are prone to change over time and in unpredictable ways. As we note in appendix D, we assumed from the start that some errors result from "deferred procedure calls" (DPCs), as they are called in the Windows domain. There may be some value in the investigation of this phenomenon. DPCs represent one type of problem under the broader class of system resource allocation issues[8] within the DAW. These have a variety of specific causes and are by definition erratic and intermittent. In addition, we remind readers about the additional errors (reported above), which seem to have resulted from other types of causes. This suggests that a first step might be the identification of a variety of causal phenomena.

  As noted, this writer foresees that this line of investigation will be resource intensive. Even if a set of causes can be identified and, later, preventive measures taken, we believe that the variable, multiple, and intermittent causes for interstitial errors rules out the development of fully reliable prevention. We believe that greater return on investment would result from the development of monitoring-and-file-management tools, as outlined below.

- **Equipment Pre-testing**

  The 2011 report noted the value, albeit limited, of pre-testing equipment to qualify it for use. For example, testing a particular brand, model and configuration of DAW to qualify it generally as acceptable for use. The sample size in our field study was too small to be conclusive but the study suggested that a given DAW may have higher than expected error rates, or have a rate higher than its peers. In other words, testing likely needs to occur for each device, and not at the brand, model and configuration level.

  Even still, testing for such error-prone behavior has traditionally been determined in non-operational, bench testing. One typical test is to run a known reference signal consisting of a pseudo-random bit sequence through the system under test and validate it on output. A properly functioning system should not alter the bit sequence. These types of analyses are incorporated into some commercial test and measurement tools, although there are no fully standardized metrics and methods.

---

8 There is an abundant literature on this topic; for one example, see http://www.informatica.si/PDF/30-1/01_Chevaleyre-Issues%20in%20Multiagent%20Resource%20Allocation.pdf

In an echo of our comments on *Cause Investigation* above, we state that periodic bench testing of a device or system with a reference signal will not replicate the environmental conditions that create interstitial errors, and we recommend instead the development of monitoring-and-file-management tools that perform on an ongoing basis while the DAW is in normal operation.

- **Monitoring the Recording Process**

There is monitoring and there is monitoring. As we noted in 2011, a variety of vendors do offer real-time (i.e. during digitization) and non-real-time (i.e. post-digitization) analysis systems. These valuable systems monitor a variety of important issues and they successfully identify many problems, including certain types of digital errors. These systems, however, generally do not detect the interstitial errors under discussion here. Interstitial errors happen during the writing of the bitstream to disk, while real-time systems monitor the bitstream prior to being written to disk. Additionally, both real-time and non-real-time analysis systems use algorithms to detect errors that are not nearly as accurate for detecting interstitial errors as the comparative analysis methods used in this study.

The desired interstitial-error-detecting, process-monitoring tool does not currently exist. This writer, however, believes that development of such a tool would be straightforward. In its simplest form it would entail building a system that records a second file, like the DDR file in the field study, with accompanying software to automate the execution of the null test. When the null test reveals an error, the system operator (or perhaps an automated agent) would look at the result. If it is clear that one file is missing something compared to the other, the good file could be retained. If the difference information is more complex or harder to interpret, the digital transfer would be repeated.

As noted in the March 2011 report, such a tool might be developed by a commercial manufacturer or by an open-source developer, especially with grant-funded support. A follow-up activity, later in 2011, delivered on these next steps generating requirements, wireframes, a workplan, and a level of effort estimate for development of an interstitial error analysis and reporting tool, and conducting a survey.[7] The field study described in this document was in a sense a sidebar, undertaken to further confirm the existence and extent of the problem and to experiment with one method for assessing these errors.

[7] The 2011 document that described the proposed tool-building project has been included in the *Appendixes* that comprise Volume II of this report. The relevant section is titled *Interstitial Error Detection & Reporting Research and Development Project Work Plan*, found on the specially numbered appendix pages 46-57.

Community Support and Action Welcome

We have been pleased to see some community interest in this topic, beyond the members of the federal agencies Working Group. We have heard from a commercial mastering studio, a preservation service bureau, and a university, all of whom have experienced interstitial errors in their work. One of these organizations has begun developing software to help automate the work of performing a null test. This writer and, I am certain, the FADGI Audio-Visual Working Group, will welcome additional commentary, analysis, and tool-building support from the recorded sound preservation community.