IRENE audio preservation at the Northeast Document Conservation Center: Developing workflows and standards for preservation projects that use innovative technology

Received (in revised form): 21st December, 2020

Julia Hawkins

IRENE Audio Engineer, Northeast Document Conservation Center, USA

Julia Hawkins is an IRENE audio engineer at the Northeast Document Conservation Center, where she manages digitisation projects, reformats historical audio formats using the IRENE imaging system and creates documentation of the IRENE technology. She previously worked as a special projects assistant at the University of New Hampshire's Diamond Library and in various positions in the library and archives of Saint Anselm College, where she earned her BA. She earned her MLIS with a concentration in archives from the University of Wisconsin, Milwaukee.

Northeast Document Conservation Center, 100 Brickstone Square, Andover, MA 01810-1494, USA
Tel: +1 978 470 1010; E-mail: jhawkins@nedcc.org

Bryce Roe

Director of Audio Preservation Services, Northeast Document Conservation Center, USA

Bryce Roe is the Director of Audio Preservation Services at the Northeast Document Conservation Center, where she confers with collection-holding institutions to evaluate their audio collections and develop preservation proposals, and manages preservation projects that use traditional technologies for magnetic and digital tape media, and either traditional or optical-scanning methods for grooved media. Bryce earned an MLIS in archives management from Simmons College, and a bachelor of arts with a concentration in ethnomusicology from Oberlin College.

Northeast Document Conservation Center, 100 Brickstone Square, Andover, MA 01810-1494, USA
Tel: +1 978 470 1010; E-mail: broe@nedcc.org

Abstract In 2014, the Northeast Document Conservation Center (NEDCC) began offering the IRENE audio preservation service to libraries, archives and museums. IRENE is an innovative optical scanning technology for digitising grooved audio carriers without using a stylus. Developed at the Lawrence Berkeley National Laboratory and the Library of Congress, IRENE uses cameras and microscopes to image the grooves at high resolution, and customised software that mimics the motion of the stylus through the grooves in the images to produce an audio file. The multi-step process offers the operator a critical degree of control for addressing the unique characteristics of 'irregular' (ie not produced in a controlled, professional environment) grooved recordings often found in archival collections, but also complicates the task of properly documenting the digital provenance of the files it produces. The rapid development of this technology presents the challenge of continually updating workflows and methods, at a rate faster than audio preservation standards and best practices can dictate. This paper provides a case study in developing preservation workflows and standards for projects that use innovative technology, by working from a foundation of established standards and providing transparent documentation of the ways in which IRENE deviates from those standards.

KEYWORDS: audio preservation, digital provenance, digital preservation, workflows, digitisation, standards

INTRODUCTION

IRENE is an innovative optical scanning technology for digitising grooved audio carriers without using a stylus playback method. IRENE is named for the first sound retrieved by the system during its development, a recording of 'Goodnight, Irene' performed by the Weavers. As an acronym, IRENE stands for 'Image, Reconstruct, Erase Noise, Etc'. The technology is very sophisticated, but the basic approach is quite simple. Essentially, IRENE uses either a two-dimensional (2D) camera or a three-dimensional (3D) confocal microscope to photograph the grooves on audio carriers like wax cylinders and lacquer instantaneous discs in ultra-high resolution. The resulting image files are then processed using custom software that translates them into an audio file.

Compared with a traditional playback digitisation workflow, whereby the analogue signal is captured, converted to digital, and encoded as a readable file in an automated single-step process while the carrier is playing, IRENE requires a multi-step imaging and analysis process to produce an audio file, see Figure 1.

Why go through the extra step of imaging the groove rather than playing it with a stylus? The principal reasons are as follows:

- There is no mechanical contact between the stylus and the grooves, which is

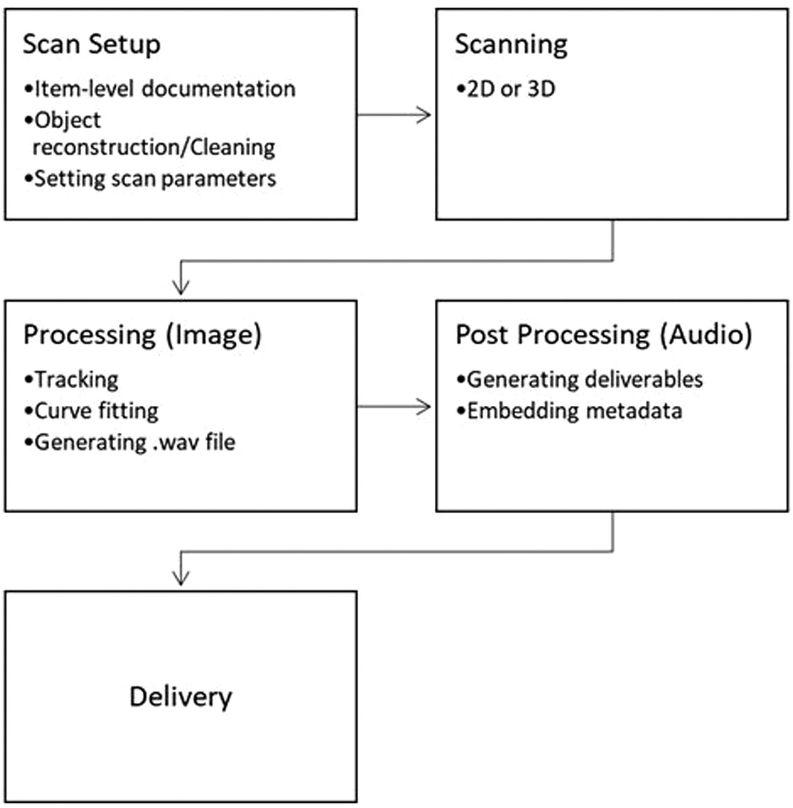

Figure 1: IRENE workflow

beneficial for fragile media, and critical for damaged or imperilled formats that cannot be played without causing further damage.

- IRENE is format-agnostic when it comes to grooved media. There is no need to acquire legacy playback equipment to digitise unusual formats.
- IRENE produces an image data file, which serves as the primary 'preservation master file' in that it captures the current condition and groove structure of the carrier and digitally 'locks it in time'.

With six systems currently in place, the IRENE technology has been used to preserve a significant number of valuable, previously inaccessible audio recordings from some of the most at-risk carriers in archival collections, such as wax cylinders, lacquer instantaneous discs, and the variety of early, experimental grooved formats. The multi-step IRENE process also presents the operator with a multitude of opportunities to consider and determine — often in a precise manner — how and what is preserved of the carrier and the recorded signal. Instead of operating as a player, IRENE offers a method by which to migrate a physical object into the digital domain, at which point a variety of large-scale data analysis tools can be applied to extract recorded sound.

IRENE AT NEDCC

The IRENE audio preservation service at the Northeast Document Conservation Center (NEDCC) is the culmination of over a decade of research and development at the Lawrence Berkeley National Laboratory (LBNL) and the Library of Congress (LOC). Development began in 2000 at LBNL by Dr Carl Haber and Dr Earl Cornell. Dr Haber learned that audio on early grooved formats was inaccessible if the format was too damaged to play, and he had the idea of applying his imaging expertise as a particle physicist to optically digitise a recording without touching the material. In 2006, a system was installed at the Library of Congress, where Peter Alyea worked closely with Haber and Cornell in order to prove the technology as a preservation tool.

In 2013, NEDCC was awarded an IMLS National Leadership Grant to bring the IRENE system out of the lab and into the marketplace as a sustainable audio preservation service. NEDCC is a nonprofit conservation and preservation centre that serves collection-holding institutions nationwide. NEDCC's goal in implementing IRENE is to save the nation's audio cultural heritage recorded on grooved media, while developing a business model that also supports the ongoing research and development efforts by the Berkeley Lab to sustain and improve the technology.

As the system continually develops and improves, it also presents the challenge of continually updating workflows and methods, at a rate faster than audio preservation standards and best practices can dictate. The following description and examples of the IRENE process at NEDCC offer a case study in the development of preservation workflows and standards for projects that use innovative technology and continually developing tools.

THE IRENE PROCESS

Imaging

The IRENE process begins with imaging. There are currently two imaging methods — 2D and 3D — which capture different kinds of data and produce slightly different file sets. Two-dimensional data is captured using a line-scan camera, pictured in Figure 2, which takes a magnified, high resolution photograph of the grooves. This method is primarily used for carriers with laterally modulated grooves, such as lacquer or shellac audio discs.

Three-dimensional data is captured using a confocal microscope. The confocal microscope focuses white light through a mixture of a lens and a prism. The different

Figure 2: 2D scanning system

wavelengths of light are dispersed by this lens and come into focus at different places, creating a multicoloured swarm of points. When an object's surface is passed through that region of multiple focus points, it produces strong reflections from the colour that is in focus and weak reflections from the colours that are out of focus. This creates a 3D map of the object's surface that can then be processed for audio.[1] This system, pictured in Figures 3 and 4, was primarily developed for digitising carriers with vertically modulated grooves, such as wax cylinders, but has also been very effective on extremely worn lacquer discs, uncoated aluminium discs and some more uncommon formats such as Soundscribers and Dictabelts.

Scanning a carrier can produce up to 60 TIFF image files for 2D scans and up to three TIFF images for 3D scans, as well as a variety of files that document the creation of those images. Each file set contains a .png snapshot of the front panel of the scanning software, and an .xml file containing all of the parameters used for that scan. Most scans use a laser ranging device that gathers height information on the surface of the object ahead of the camera, helping the camera to adjust its height and remain in focus. The focus data is stored in a discrete file, with a unique extension depending on the scanning method used. A host of other files may be generated, especially during 3D scans. A .pri file contains height information from the object. A .bri file contains information about the intensity of light at each measurement. Each of these files has the potential to be helpful to IRENE operators during analysis and troubleshooting, but only the TIFF file(s) will ultimately be processed.

PROCESSING

Images are processed using a specialised program called Weaver, developed and maintained by Dr Earl Cornell at LBNL. Images are processed using a series of plugins, each with their own discrete function. This modular processing software measures the lateral or vertical displacement of the groove, with respect to its unmodulated trajectory, along a sequence of points (pixels). The stylus velocity at each point along this displacement waveform is measured using numerical differentiation. With this information, it is able to reproduce the sound recorded on an object in the form of a .wav file.[2] When an IRENE operator is processing an image, they are primarily focused on setting parameters within Weaver to perform two functions: tracking and curve fitting.

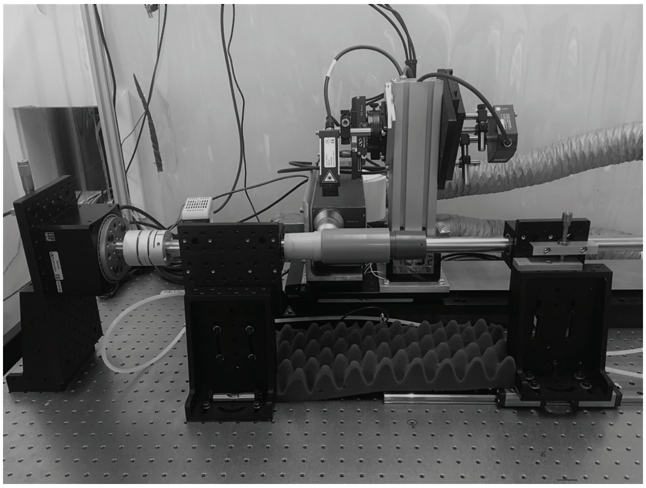

Figure 3: 3D scanning system

Figure 4: Confocal microscope scanning a wax cylinder

On an undamaged carrier (pictured here in a 2D image), tracking usually involves some combination of binning and smoothing an image to produce a continuous black line along which a track will be placed (see Figure 5).

The track is then generated automatically based on the image data provided and the parameters set in the tracking plugin. It should ideally correspond to the position of the groove bottom in the original image data (see Figure 6).

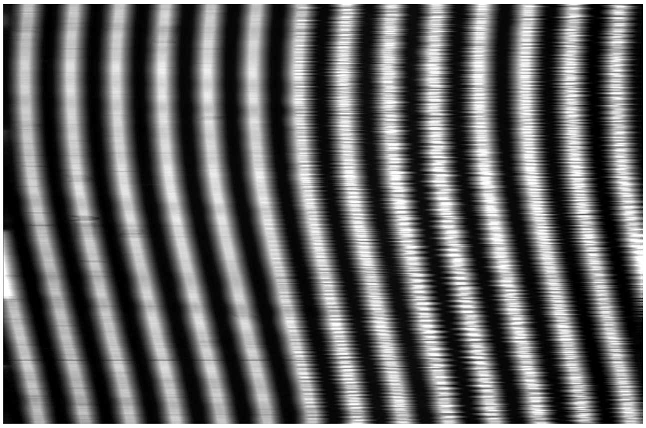

Figure 5: 2D image ready to be tracked

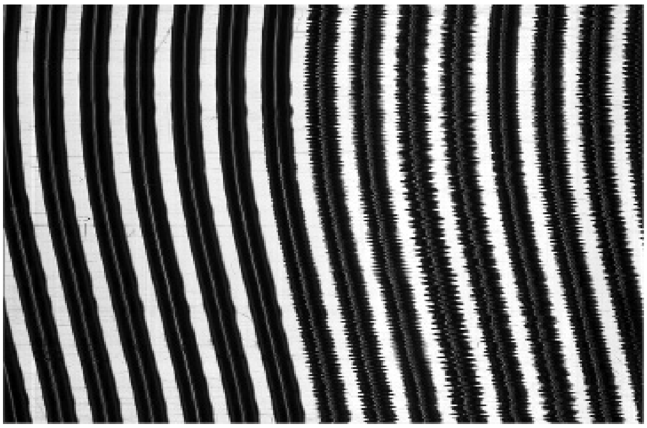

Figure 6: 2D image tracked

When processing damaged media, creating a continuous track is usually more labour-intensive. Although the tracking parameters can be adjusted to automatically accommodate some discontinuities, the path of the groove may be severely disrupted in a variety of ways that cannot be accurately tracked without the intervention of the operator. Carriers such as wax cylinders, shellac discs and glass-based lacquer discs are fragile and may be broken in multiple pieces. Lacquer discs are also susceptible to delamination. As the discs age, the soft lacquer will react differently to environmental conditions than the rigid base material (usually metal or glass). This usually results in a loss of plasticiser in the lacquer, producing exudation of palmitic and stearic acid on the disc surface.[3] The exact nature of this deterioration will vary from disc to disc, as shown in Figures 7 and 8, but often the lacquer will shrink slightly, crack, and eventually lift away from the base material of the disc, leaving modulated pieces of lacquer that must be puzzled together physically and virtually.

The advantage of using a modular program like Weaver is that the image data from a scan can be manipulated to address these different types of damage 'in post'. Weaver currently supports almost 200 different plugins, thus providing a wide variety of options to explore when trying to extract the best possible audio from a worn or damaged carrier. In addition, plugins can be written or updated to address challenges as they arise.

Figure 7: Lacquer lifting from aluminium base in irregular strips

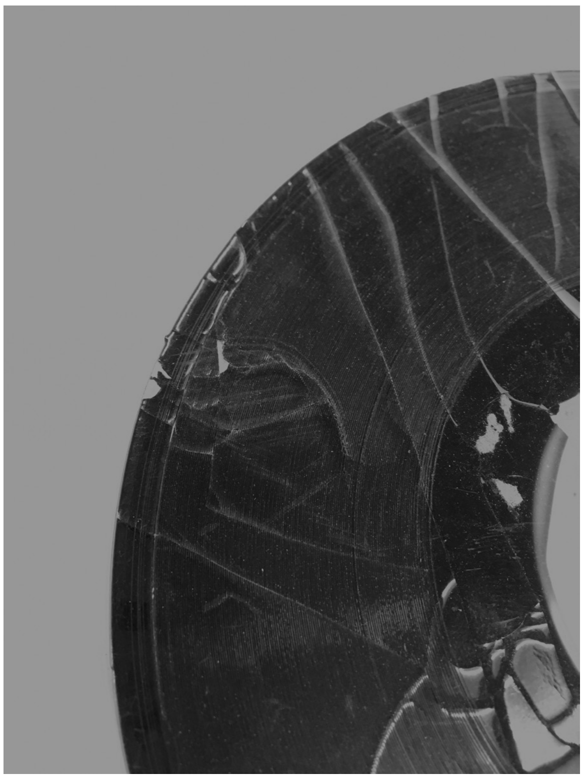

Figure 8: Ridged delamination

If a carrier is broken, every effort is made to physically align the pieces before scanning. However, when analysing something on such a small scale, the slight shifting that can occur when aligning multiple pieces of a broken object often results in misalignment that the tracking algorithm cannot handle automatically. In these cases, automatic tracking would produce disjointed files that jump around in the recorded programme. To avoid this, the operator could use the plugin MoveRegion to select an area of the image and move it until the grooves are properly aligned (see Figure 9).

In some cases, the carrier may be cracked or delaminating in large, irregular sections that MoveRegion cannot easily address. Instead of manipulating the image data, an operator could use the plugin ManualCracks, which allows the operator to drop data points along a crack and indicate the direction in which the tracking algorithm should try to 'jump' when it crosses the break. This increases the chance that as it runs, it will place the track in the right location without constant correction by the operator.

If a crack is too large, the white space of the crack may cause the tracking algorithm — which is effectively trying to follow a continuous black line — to 'lose its place' as it crosses the crack. To correct for this, the operator could use the plugin FillRegion to drop data points along the edges of the crack and fill the white area with a stretched version of the image data above and below the empty space (see Figure 10). This does not replace lost audio, but rather gives the tracking algorithm something to follow across these gaps.

Even the help of these tools may not be enough to facilitate automatic tracking on some images. IRENE is often the only option for digitising extremely damaged and delaminating carriers, and in these cases operators will manually create a track through the severely disrupted surface (see Figure 11).

Previously, this entailed dropping individual points along the groove, but IRENE developers and operators have recently shifted to using a hybrid model. The tracking plugin attempts to create a track automatically as it runs, but it can be paused and edited along the way. This development has drastically decreased the processing time for extremely damaged carriers. It has also helped address persistent issues like properly tracking over-modulated areas (see Figure 12).

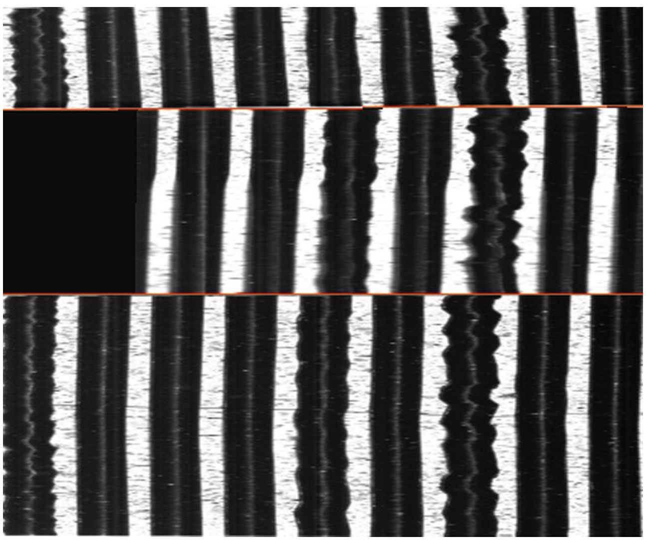

Figure 9: Aligning grooves using MoveRegion plugin

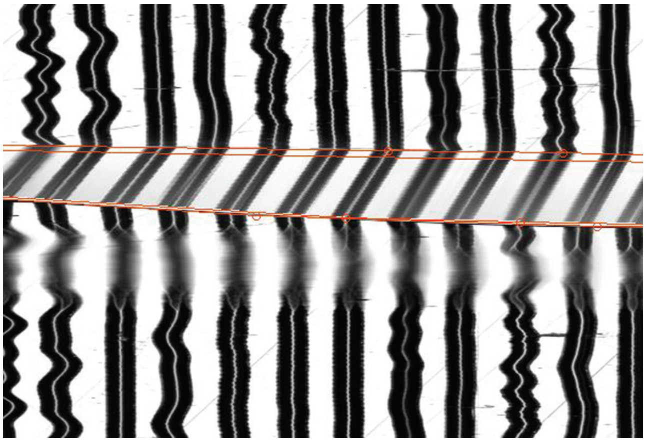

Figure 10: Filling cracks for tracking with FillRegion plugin

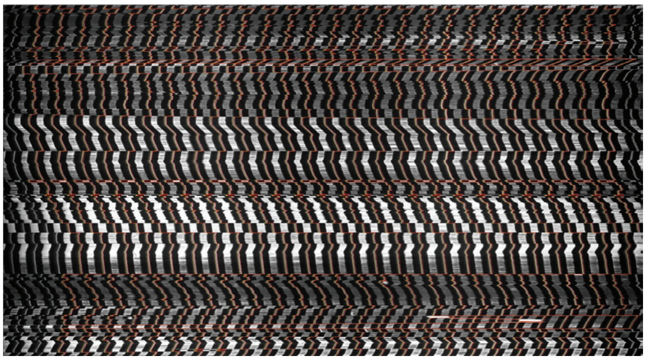

Figure 11: Two manually tracked TIFs from a badly delaminating lacquer disc

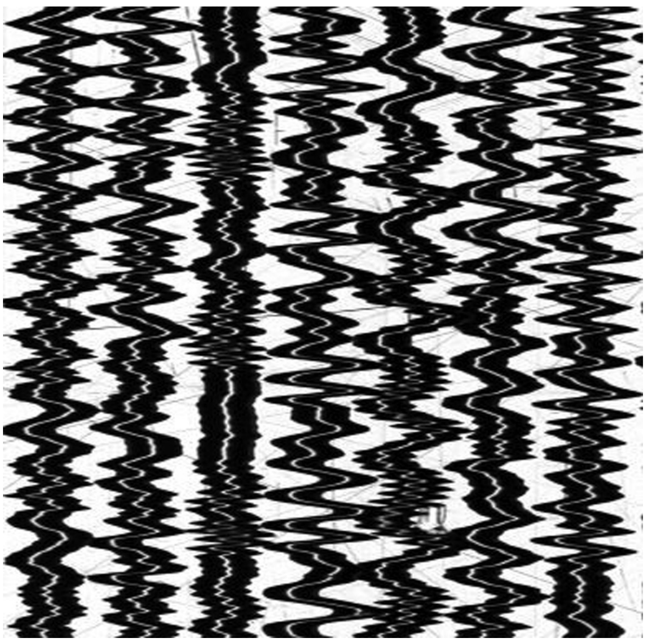

Figure 12: Severe over-modulation

Manual tracking can also be successfully used to track areas where one recording is entangled in or overlaid by another. Sometimes those recording a programme attempted to record a 'second pass' in the spaces between the first set of grooves. Sometimes wax cylinders were not shaved down enough to fully eliminate all traces of an initial recording before they were reused. Manually controlling track placement gives IRENE operators the option to recover these recordings in part or completely depending on the characteristics of the recording.

Once a track is generated, either automatically or manually, Weaver runs a plugin that follows that track and creates a 'fit' along the curve of the groove. The accuracy of that fit is determined by the accuracy of the track and the parameters set within the curve fitting plugin by the operator. The data from the fit and the track will be used to create the resulting .wav audio file.

CONTROL AND SUBJECTIVITY

Generating audio with IRENE is a series of subjective decisions on the part of the operators. To some extent, the decisions made during scanning and processing are based on the characteristics of the original carrier. For example, the optical sample rate of a scan is calculated based on the original recording speed of the carrier, and the desired sample rate of the audio. Other scanning parameters, such as the exposure and intensity, can fluctuate based on the condition of the object, whether these settings interact with focus data at certain thresholds, and to some extent on the experience and tendencies of the operator. Processing involves an even greater amount of subjectivity. The average plugin chain can contain anywhere from 10 to 20 individual plugins, which can each have up to ten adjustable parameters. There are tools in the software and visual cues on the display that help an operator determine how these parameters should be set based on the data at hand, but there is still a great degree of flexibility. The operator is usually working with a basic plugin chain used on similar images, and makes adjustments based on what produces the 'best' audio from that image.

This approach offers a critical degree of control for addressing the unique characteristics of compromised and/or 'irregular' (ie not produced in a controlled, professional environment) recordings often found in archival collections. It facilitates creative problem-solving without risking the physical object at hand. It also allows IRENE developers and operators to continually update their scanning and processing tools and strategies to address new challenges. With greater control, however, comes the responsibility to document the choices made by operators as completely as possible.

The role of the archivist as subjective arranger and describer has been discussed extensively in the professional literature of the last few decades. In her paper 'Picking our text: archival description, authenticity, and the archivist as editor', Heather MacNeil examines this ongoing discussion and the relationship between transparency in archival description and the authenticity of historical records. She states:

> archival description involves conscious and deliberate decisions about the representation of archival documents. And because description constitutes the frame of reference that shapes the meaning and significance of those documents, archivists are obliged to render an account of our role and responsibility in the process of our representation.[4]

Although IRENE operators are not the archivists responsible for the long-term preservation of the files they generate, by actively participating in the digitisation of historical audio recordings, they effectively create new 'representations' of those recordings, becoming co-creators of those records. The aim of this process is always to generate a digital file that accurately replicates the original recording, but the IRENE system simply allows the operator to make more choices than might be afforded using a different method of digitisation. To guide those choices, the IRENE lab at NEDCC has developed its workflows and frameworks for file creation by looking to established standards for the digital preservation of audio.

PRESERVATION ETHICS AND STANDARDS AT NEDCC

The lifespan of audio carriers is limited by their physical and chemical instability, the availability of playback equipment, and the fact that the original equipment may be a potential source of damage. Consequently, audio preservation relies on the production of duplicates that can stand for the original as 'preservation surrogates'. The International Association of Sound and Audiovisual Archives (IASA) Technical Committee, 'Guidelines on the Production and Preservation of Digital Audio Objects' is the accepted authoritative guide on how to create digital audio preservation surrogates, including signal extraction from original carriers, and preservation target file formats and specifications.[5] While the IASA guidelines were developed around traditional playback methods for reformatting, IRENE workflows at NEDCC are designed to meet the same established preservation goal: to create a faithful and accurate preservation surrogate of audio archival material in its current condition, in a format that can be carried into the future, ensuring access to the audio content for current and future users.

As stated previously, in a traditional playback digitisation workflow, the analogue signal is captured, converted to digital, and encoded as a readable file in an automated single-step process while the carrier is playing. Therefore, the appropriate selection and calibration of each piece of equipment in the signal chain is critically important in the ability to create a digital copy with high precision and integrity. The analogue-to-digital convertor (A/D) is responsible for transforming the signal from its analogue source to a digital approximation, and is therefore considered the most critical component in a playback digital preservation pathway for accuracy. In converting the analogue signal, IASA notes that the A/D 'should not colour the audio or add any extra noise'.[6]

Because IRENE was conceived and developed as a preservation tool, its functionalities are designed to enable the operator to take preservation actions towards creating a faithful, accurate digital surrogate, without adding anything extra. For example, deterioration of audio carriers due to age, mould or other damage, are irreversible, and full digital audio restoration is beyond the scope of a preservation-oriented service. With IRENE, the software functionally enables stabilisation of damaged media to enable accurate retrieval of the remaining signal (eg aligning the grooves in the image for processing, where damage may have caused misalignment of grooves), but where signal is lost (eg a gauge in the grooves of a cylinder, or a separated and lost section of lacquer on a disc), the processing software does not function to extrapolate, recreate or fully restore lost signal.

As with any digital object created for preservation purposes, embedded metadata is essential for facilitating long-term management, preservation and interpretation of the contents.[7] The Federal Agencies Digital Guidelines Initiative (FADGI) Audio-Visual working group articulates common and sustainable technical guidelines, methods and practices for digitised and born-digital sound recording and moving images, and NEDCC follows FADGI's 'Guidelines: Embedding Metadata in Broadcast WAVE Files' to the extent that they sufficiently support essential preservation documentation of the IRENE files. For example, in a traditional workflow, much of the metadata about the digital provenance of the file, or how the digital file was created, is revealed by listing the exact equipment used in the signal path. FADGI recommends recording these important details in a 'coding history' field to 'describe the signal chain from which the digital file was created, starting with the analogue or digital source'.[8] FADGI provides strict guidelines for formatting this field, but the free text element allows for some flexibility. The current template for an embedded IRENE coding history describes the original carrier, indicates that it was scanned with IRENE, identifies Weaver processing and the audio program used to create the final version of the audio at the standard sample rate:

A=ANALOGUE, M=mono,
  T=IRENE3D; LacquerDisc; 78RPM
A=PCM, W=32, M=dual-mono,
  T=Weaver; Asus; Windows7
A=PCM, F=96000, W=32, M=dual-mono,
  T=iZotopeRX5; Asus; Windows7

FADGI's formatting guidelines for this field do not enable the inclusion of the multitude of lighting, scanning and processing decisions and actions that have occurred to create the digital file. As the IRENE process involves additional decision points and preservation actions, where IASA and FADGI's instructions for producing digital audio surrogates with appropriate metadata fall short, NEDCC has developed workflows and deliverables that meet the preservation objectives they define.

DELIVERABLES

NEDCC's standard set of IRENE deliverables has developed with the technology and with an understanding of what institutions want and need in order to best preserve these file sets over time. While this standard incorporates the ethics and standards identified above, it does deviate from 'best practice' where there is demonstrable benefit to the quality and completeness of the files. It will likely continue to evolve as the community of IRENE operators grows and conversation continues with the preservation community at large. Although IRENE operators will generate files according to client specifications if they are provided, at present, the following is a complete inventory of the recommended IRENE deliverables.

First, the audio preservation master file is provided in the uncompressed Broadcast Wave Format (BWF). This nonproprietary format is commonly viewed as the standard for archival audio, as it retains the widely recognised .wav extension of its predecessor, while supporting a wider range of embedded metadata, namely the 'Broadcast Audio Extension' (BEXT) chunk.[9] In accordance with FADGI guidelines, IRENE master files are embedded with an MD5 checksum and the aforementioned coding history.

Per IASA guidelines, preservation master files should have a sample rate of at least 48 kHz,[10] but for most projects, master files are provided at 96 kHz, with the exception of wax cylinder projects, which are delivered at 48 kHz unless otherwise requested.[11] Master files should also have a bit depth of at least 24,[12] but the IRENE lab at NEDCC has recently shifted to delivering 32-bit floating-point audio master files.

To understand how high the resolution of these master files actually is, and why master files are delivered at 32-bit floating-point rather than fixed-point, it is helpful to understand some basics of digital audio. Sound waves are compressions and rarefactions in the air. Recorded sound of any kind captures these compressions and rarefactions over time. Storing sound digitally requires breaking the recorded information up into discrete data points that each represent some amplitude at a particular point in time. This process is called quantisation, where the data points, called samples, are assigned a binary value consisting of 16 bits, 24 bits, or 32 bits.

Sample rate determines how many data points are gathered per second (Hertz). The sample rate determines the highest frequency that can be accurately represented. The range of human hearing is roughly 20 Hz to 20 kHz. In order to accurately capture the highest desired frequency, one must adhere to the Nyquist theorem, which states that the sample rate of a recording must be double that of the highest desired frequency.[13] This combats a digitisation error called aliasing, which can result in audible artefacts. For example, in order to accurately capture the upper limit of human hearing (20 kHz), according to Nyquist, the sample rate must be 40 kHz. The sample rate for IRENE preservation master files is 96 kHz which, again, has an upper frequency limit of 48 kHz.

Bit depth, also called word length, determines the accuracy of the data points that are sampled, and ultimately the dynamic range of the recording. It does this by determining the number of possible values that can be represented by a single sample. The more possible values available, the more accurate that sample will be. A 16-bit fixed-point file has 16 ones and zeros that are used to describe each sample. 16-bit is the standard for CD quality audio. A 16-bit word can represent 65,536 discrete values. A 24-bit file is sampled with 16,777,216 discrete values.[14]

The discussion of bit depth so far has dealt with fixed-point arithmetic. Again, 16-bit fixed-point assigns 16 ones and zeros to each sample. These binary bits represent a specific value and no further calculations are needed to generate its value. Floating-point arithmetic uses a different method to encode and decode binary information. The main difference is that floating-point arithmetic uses scientific notation ($a \times 10^b$) to describe each sample. Because of the use of scientific notation, floating-point arithmetic allows for many more potential values, compared with fixed-point.[15]

The Weaver software automatically generates files at 32-bit floating-point, and when examining the way in which the data is affected, it is easy to see why. While 24-bit fixed-point files can capture audio signals with a great degree of accuracy, 32-bit floating-point files offer even more accuracy and flexibility. For floating-point audio files, the possible number of values per sample is much higher than fixed-point, enabling the processing software to perform more accurate calculations, and to record and reproduce audio information that would have been distorted at a lower bit depth. Due to the nature of historical audio recording, IRENE files may contain content that was originally recorded at a very low volume, overlaid with extremely high-amplitude noise spikes from damage and wear. In a fixed-point file, these large noise spikes will distort or 'clip' wherever the recorded sound's amplitude exceeds the maximum value of the fixed-point word. In a traditional analogue-to-digital transfer process, a fixed-point master file might be recorded at a slightly lower gain to ensure that the signal remains comfortably within the usable dynamic range of the fixed-point recording. For IRENE files, however, this would mean drastically adjusting the scaling factor for the gain in the plugin that generates the .wav file, creating IRENE master files with nearly inaudible content.

For IRENE, 32-bit floating-point files are delivered because they can accurately describe audio with these great ranges in dynamics, where clipping distortion is barely an issue. It allows for delivery of unmanipulated master files where, in this example, the quiet audio content can be presented at its original level and maintain the integrity of any high-amplitude noise spikes, which are still a part of the recording and represent the audio from the item in its current condition. In other words, content that 'clips' in 24-bit fixed-point files does not 'clip' in 32-bit floating-point.[16] The significantly increased 'headroom' of 32-bit floating-point makes it possible to capture the full dynamic range of the recording, without creating files that obscure already difficult to understand content.

The IRENE operators at NEDCC have also shifted to delivering master files for laterally modulated materials as 'dual-mono' files, even though most of those recordings were originally mono.[17] These two channel files are the audio generated from the right and left side of the groove. On a perfectly preserved mono recording, both these files would be the same. On many worn carriers, however, this is not the case. The stylus used for playback on some carriers may have been misaligned and worn away one side of the groove more than the other. Some carriers may exhibit crazing or palmitic acid exudation that will obscure audio intermittently. In some cases, delaminating lacquer discs might have a tear along the groove bottom, as shown in Figure 13, that leaves only one side of the groove sufficiently intact to be analysed.

Analysing image data allows the operator to see these anomalies in great detail, and make an informed decision about how to process the image to produce the best audio. Unfortunately, it is very rare for one side of the grooves to be consistently better than the other, so where there is variation, it is captured in the generated master file.

Another perk of analysing image data is the ability to optically 'clean' files. When fitting the curve of the groove during processing, outliers in the data may be incorporated into the fit. This may be caused by a piece of dust, wear on the groove wall, or other factors, and will cause a pop or click in the audio. During curve fitting, parameters are set to limit these outliers as much as possible, but some additional tools help to make these clicks and pops even less noticeable in the final audio. For laterally modulated materials, the plugin CleanTrack helps to 'trim' excess noise from a file, essentially creating a high-frequency roll-off.

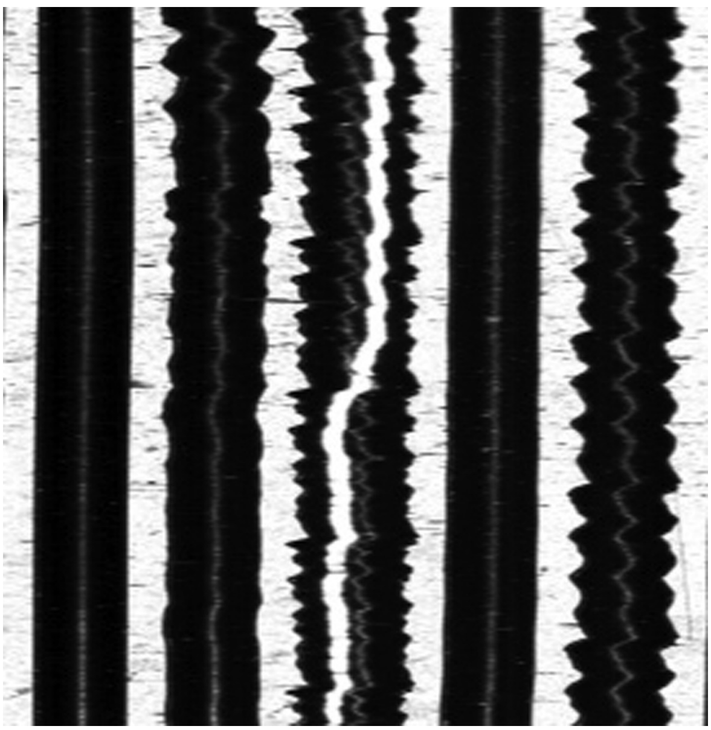

Figure 13: Tear along the groove bottom

The corresponding process for wax cylinders, which are susceptible to mould damage, is called 'BlobClean', which takes advantage of the height information gathered in 3D scans to differentiate between surface-level mould and modulation of the grooves. These areas are highlighted in a 'mask' of the image, shown in Figure 14, and then excluded from the data used during curve fitting. This does not replace the signal lost from the mould damage, but it does eliminate the noise that would be created by the mould, resulting in files that are a little more listenable.

In most cases where this optical cleaning is applied, NEDCC delivers the 'raw' file produced before optical cleaning, and generally uses the 'clean' file to generate the master file. The 'raw' version then serves as a digital surrogate that presents all options for creation of alternatively processed files.

In addition to the audio preservation master file, NEDCC also provides two access files.[18] The first is a higher-resolution 'intermediate' 48 kHz, 16-bit file, and the second is a direct MP3 rendering of the intermediate file, at 160 Kbps per channel. These files are processed using commercial audio editing software that will improve the listening experience, or in some cases aid in bare bones comprehension of the audio on carriers that have suffered extreme damage. These techniques include de-clicking, de-noising, equalisation, peak-normalisation, etc. The goal is not to provide a perfectly remastered copy of the audio, but to make the content of the audio immediately accessible.

At this point, specific guidance from the audio preservation community ends. This is not a failure of the field, but a natural consequence of working with new technologies. Until IRENE becomes a more accessible and widely implemented technology, many of the files generated by IRENE will be somewhat unfamiliar even to preservation professionals. It would be irresponsible to focus on providing access to the content on these at-risk recordings while obscuring the process that has made them accessible. On the other hand, in both the imaging and processing stage, IRENE generates a large number of files, many of which have no purpose beyond initial analysis and potential troubleshooting. Archivists and collection managers have practical concerns such as storage space and the long-term preservation of files which must be taken into account. It would be impractical to hand off bulky sets of files to institutions, when (1) many of the files would be inaccessible without the

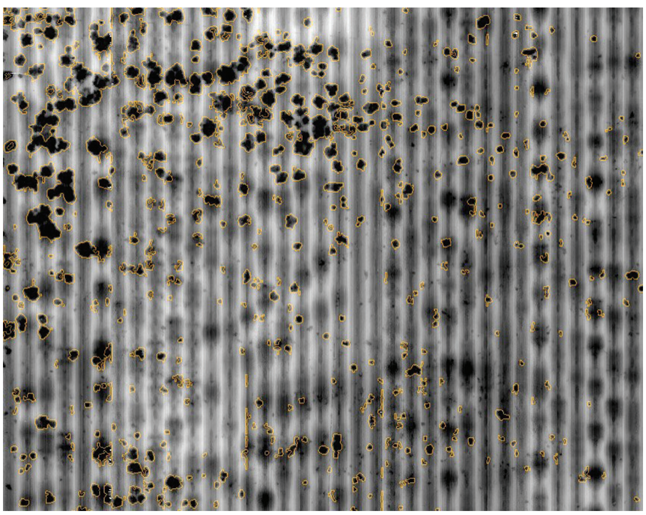

Figure 14: Mould identified by BlobClean

appropriate software, and (2) many of the files were not used in any substantial way during the creation of the audio deliverables.

It is the responsibility of IRENE operators to provide information to the individuals and institutions responsible for the management of these files so that they can make informed decisions about their long-term preservation, description and accessibility. As previously discussed, the audio files generated by IRENE, even with standard embedded metadata, do not provide this information when removed from the context in which they were created. With this in mind, NEDCC provides the following files to accompany every audio file set:

- *The images from the IRENE scans:* These serve as the 'image preservation master file(s)'. These are written in the widely-accepted, fully documented TIFF file format, which can be opened in a variety of common software applications. It is the preferred format of the Library of Congress for 2D and 3D bitmapped images and accompanying image files for audio.[19] These images become a surrogate for the original carrier, in that they capture the current condition and groove structure of the carrier and digitally 'lock it in time' as the physical carrier continues to degrade. As the IRENE technology continues to develop, these images could potentially be reprocessed to produce superior audio in the future. In addition, these images provide their own source of data about the original carrier that may not otherwise be accessible to researchers.
- *The item-level report:* This is a .xlsx file containing information on each original carrier, and relevant metadata from the original object, imaging and audio processing.
- *The summary report:* This serves to contextualise this information by providing a description of the entire IRENE transfer process, highlighting any issues encountered during the project and how they were addressed, and providing a descriptive inventory of the deliverables.
- *The .xml file generated during scanning:* This contains a record of all parameters used during a particular scan.
- *The .plg file generated during processing:* This is a formatted text file that contains the complete plugin chain and specific parameters used to create a given audio file. These files provide an easily accessible source of metadata that can be opened in Weaver, or in any basic text editor.
- *All data files created during processing (.trk, .blb, .crk, etc):* Every time an IRENE operator performs a manipulation on an image in Weaver, the relevant data points used to do so are saved in a data file linked to that specific plugin and its position in the plugin chain. In addition, any tracking that is not completely automatic generates a track file (.trk). These are .xml files that contain *x* and *y* coordinates for every point in a manually generated track. Each of these data files is relevant to a specific image and is often labour-intensive to create. For those who do not work with Weaver, these files may appear useless, but they would be crucial to replicating the initial processing of any IRENE images and potentially improving upon that processing in the future as the technology continues to improve.

These files have been identified as essential to maintaining the digital provenance of IRENE audio files. Although these files represent a fraction of the data generated during the IRENE digitisation process, together they contain a full and contextualised report of the decision-making process that produced an IRENE audio file.

CONCLUSION

The preservation of early audio formats is a race against time. Digitisation with IRENE is not the perfect solution for every audio preservation project, but it is an incredibly valuable tool for preserving some of the most vulnerable pieces of recorded history. As the number of systems in use around the world continues to grow, so will the ability to share experiences and reach consensus on standards and best practices. But entropy will not wait for that consensus. The audio preservation community has provided an ample foundation on which IRENE operators can base their decisions today. IRENE was developed as a preservation tool, and the community of scientists, archivists, audio engineers and students that have contributed to its development are guided by the same goals: to make at-risk audio collections accessible and to provide the documentation and transparency that will support their long-term preservation. It is the hope of IRENE operators at NEDCC that further discussion of the intersections between established standards and IRENE workflows will continue to improve the best practices of IRENE operators around the world. In addition, the IRENE operators at NEDCC hope that the methods employed here for developing workflows around new and rapidly changing technologies may be useful to other preservation professionals working outside the scope of established standards to preserve at-risk media.

References

1. Fadeyev, V. Haber, C., Maul, C., McBride, J.W. and Golden, M. (2004) 'Reconstruction of recorded sound from an Edison cylinder using three-dimensional non-contact optical surface metrology', *Journal of the Audio Engineering Society*, Vol. 53, pp. 264–278.
2. *Ibid.*
3. Brylawski, S., Lerman, M., Pike, R. and Smith, K. (eds.). (2015) 'ARSC Guide to Audio Preservation', Association for Recorded Sound Collections, Eugene, OR.
4. MacNeil, H. (2005) 'Picking our text: Archival description, authenticity, and the archivist as editor', *The American Archivist*, Vol. 68, No. 2, pp. 264–278.
5. International Association of Sound and Audiovisual Archives Technical Committee. (2009) 'Guidelines on the Production and Preservation of Digital Audio Objects', 2nd edn, available at: https://www.iasa-web.org/audio-preservation-tc04 (accessed 18th November, 2020).
6. *Ibid.*
7. *Ibid.,* ref. 3 above.
8. Federal Agencies Audio-Visual Working Group. (2012) 'Embedding Metadata in Digital Audio Files: Guideline for Federal Agency Use of Broadcast WAVE Files', available at: http://www.digitizationguidelines.gov/audio-visual/documents/Embed_Guideline_20120423.pdf (accessed 18th November, 2020).
9. *Ibid.*
10. *Ibid.,* ref. 5 above.
11. National Archives and Records Administration. (2020) 'AUD-P2 — Audio Limited Capture', available at: https://www.archives.gov/preservation/products/products/aud-p2.html (accessed 18th November, 2020).
12. *Ibid.,* ref. 5 above.
13. Pohlmann, K.C. (2011) 'Principles of Digital Audio', 6th edn, McGraw-Hill, New York, NY.
14. *Ibid.*
15. *Ibid.*
16. Sound Devices. (2020) '32 bit float files explained', available at: https://www.sounddevices.com/32-bit-float-files-explained/ (accessed 18th November, 2020).
17. *Ibid.,* ref. 5 above.
18. *Ibid.,* ref 3 above.
19. Library of Congress. (2020) 'Recommended Formats Statement. 2020–2021', available at: https://www.loc.gov/preservation/resources/rfs/RFS%202020-2021.pdf (accessed 19th November, 2020).