Harvard University Library
Library Digital Initiative Grant

FINAL REPORT

Music from the Archive: A New Model of Access to Rare and Unique Sound Recordings

Eda Kuhn Loeb Music Library
Harvard University

Submitted by Peter Laurence
Eda Kuhn Loeb Music Library
October 1, 2007

Table of Contents

Introduction

As part of Harvard's Library Digital Initiative, *Music from the Archive* focused both on preservation and on broadening access to the wealth of rare and unique sound recordings held in the Eda Kuhn Loeb Music Library. In this ambitious project we sought not only to produce a specific body of useful *and* durable digital resources, but also to develop a knowledge base that would serve future efforts at providing networked access to our collections. Many of the valuable recordings we selected for this project existed on fragile open reel tapes that were over 40 years old. Thus, preserving the highest quality digital copy, and creating a lasting format were primary concerns. We sought durability not only in the digital audio formats themselves, but also in our descriptive metadata systems and in any facilities we might use for storing these files.[1]

In this project we proposed to offer expanded access by creating multimedia electronic finding aids that integrated digitized audio files and other digital objects. With these finding aids we hoped to enrich the research potential of our collections, and also to begin showing the relationships between the different parts of these collections. For example, how does a researcher's notebook relate to a concert program and then to a sound recording?

In late 2000 we renovated our audio preservation studios and acquired the equipment necessary to begin the work of LDI 2. Starting from scratch we imagined an infrastructure for the creation, delivery and management of our digital resources. In fact, at that point very little was in place other than the Digital Repository Service (DRS), and no one knew very much about the processes we would eventually develop in this program. We had very basic questions about formats, delivery systems and how to use the DRS. Fundamentally we needed a reliable way to combine audio and its related metadata together so that it could be stored in the DRS.

A peripheral goal of the project was to train our own permanent staff in each newly developed process and retain the expertise as much as possible. By integrating these practices into the library's existing workflow we ultimately hoped to move from a project to a continuing program.

---

[1] Danielson, Virginia. "Stating the Obvious: Lessons Learned Attempting Access to Archival Audio Collections." Paper presented at the conference *Folk Heritage Collections in Crisis*, Library of Congress, Washington, D.C., December 1–2, 2000.

From the outset the *Music from the Archive* project was poised to make a number of unique contributions to the university's LDI efforts. First among these was the establishment of a full audio preservation program at Harvard. We also needed to find out whether the standard for electronic archival finding aids, Encoded Archival Description (EAD), could successfully be applied to audio collections. The EAD format did prove flexible enough for our purposes, though this was all new work and it took considerable time and effort. Along the way we also made a substantial contribution to the development of audio metadata standards. And finally, we created two new software tools for the collection of metadata and the deposit of audio files into the repository.

The project focused on three of the library's major collections: The Laura Boulton Collection of Byzantine and Eastern Orthodox Chant, the Joseph Jeffers Dodge Duke Ellington Collection, and the James Rubin Collection of South Indian Classical Music. The Boulton collection consists of more than 400 hours of field recordings made in Greece, Eastern Europe and the Middle East during trips in the 1960s. The Dodge collection is comprised of 2,036 commercial recordings and 44 open reel tapes by Duke Ellington and a variety of his colleagues. Finally, the Rubin collection is an unrivaled treasury of more than 2,000 live recordings by the finest Indian musicians of the century. All three collections include printed field notes, program books and other documentation as well as sound recordings. The Boulton and Rubin recordings were made on open reel tapes, and the Ellington collection consisted of primarily shellac and vinyl discs.

The scope of this project required that much of the work be done in close association with several related Harvard agencies. As we depended on the infrastructure and execution of these groups, many of whom had relatively little experience at that point either, our timeline proved to be both innovative and eccentric. We were finding our way with few standards in place, and the project understandably stretched LDI's available systems and services. A lot of the work we set out to do was simply more difficult and much more time-consuming than we had predicted. Still, there were many successes and *Music from the Archive* added substantial and unique content to the University's growing collection of digital resources. Producing nearly 9 gigabytes of real audio files, the final tally of sound files created and deposited in the DRS came to 503 archival masters, 999 delivery files and 89 production masters.

Project Staff

Project Manager:

David Ackerman, Audio Preservation Engineer, Eda Kuhn Loeb Music Library, Music Building.

Project Associates:

Sarah Adams, Keeper of the Isham Memorial Library, Eda Kuhn Loeb Music Library, Music Building.
Virginia Danielson, Richard F. French Music Librarian and Curator of the Archive of World Music, Eda Kuhn Loeb Music Library, Music Building.
Robert Dennis, Recordings Librarian, Eda Kuhn Loeb Music Library, Music Building.

Project Team:

Carolann Buff, Project Assistant, Eda Kuhn Loeb Music Library, Music Building.
John Loud, Project Cataloger, Eda Kuhn Loeb Music Library, Music Building.
Doug Freundlich, Project Assustant, Eda Kuhn Loeb Music Library, Music Building.
Brian Hoffman, Cataloger's Assistant, Eda Kuhn Loeb Music Library, Music Building.
Bruce Gordon, Audio Engineer, Eda Kuhn Loeb Music Library, Music Building.

Project Advisors:

Wendy Gogel, Digital Projects Program Librarian, Office for Information Systems, Harvard University Library
MacKenzie Smith, Digital Library Projects Manager, Office for Information Systems, Harvard University Library
Robin Wendler, Metadata Analyst, Office for Information Systems, Harvard University Library

Lee Mandell. Programming Analyst, Office for Information Systems, Harvard University Library

Related Harvard Agencies and Services:

Office of Information Systems (OIS)
Digital Repository Service (DRS)
Harvard College Library Digital Imaging Group (HCLDIG or DIG)
University Information Systems (UIS)
Information Technology Services (ITS)
Name Resolution Service (NRS)
Page Delivery Service (PDS)
Streaming Delivery Service (SDS)

Staff Summary:

Over the course of the LDI 2 project we utilized existing staff members, created several new positions, and in a few cases outsourced specialized programming tasks. We also worked closely with and relied on several Harvard agencies.

- Dave Ackerman was project manager & audio engineer.
- Sarah Adams was project associate & finding aids manager.
- Virginia Danielson was project associate, overseeing finances and intellectual property issues.
- Carolann Buff was multifaceted project assistant.
- John Loud was subject specialist & cataloger for the Rubin collection.
- Bruce Gordon was our second audio engineer, completing the final deposit work of the *Music from the Archive* grant.

Note: We gratefully acknowledge the contributions made by colleagues at the Harvard College Library Digital Imaging Group, the Digital Repository Service, and most of all the Office of Information Systems. Special thanks go to Mackenzie Smith, formerly of OIS.

Methods/Processes

The work of LDI 2 can be divided into these basic processes:

- **Finding aids**: inventory and mark-up; encode links to sound files and supporting documents; upload to OASIS
- **Supporting materials:** submit test scans; create metadata instructions for each object; scan at DIG; deposit images into DRS; link images to finding aids
- **Reformatting:** create high density AIFF *archival* file; from that derive a CDR**;** from that derive a mid-density AIFF; from that derive a Real Audio file
- **Deposit of sound files:** deposit AIFF and Real Audio files into DRS; name files in NRS; link files to finding aids

The underlying goal of *Music from the Archive* was to define and develop an audio submission information package (SIP) that could be used for this and future projects. This SIP contained the original audio and all metadata in a format suitable for its deposit into the DRS. During the course of the project two important software tools were also created to support our work. The first, a core audio or audio object tool, facilitated the deposit of reformatted audio files into the DRS; the second, a process history tool was designed to capture metadata on the work history of each sound file. The final deposit of material for all three collections had to wait for the completion and testing of these tools.

Project Timeline

Selection

April 2000 – Curators of Boulton and Ellington collections selected audio recordings to be digitized.

Permissions

July 2000 – All legal permissions for recordings to be digitized had been acquired.

Finding Aids

Boulton

October 2000 – Boulton finding aid completed and loaded into OASIS.
August 2002 – Uploaded images to finding aid.
February 2005 – Audio files deposited and linked to the finding aid.

Ellington

July 2000 – Inventory of materials completed.
October 2000 – Finished markup of finding aid.
August 2002 – Finding aid completed and loaded into OASIS.
February 2005 – Deposited audio files linked to the finding aid.

Rubin

June 2001 – Completed inventory; began assembling the finding aid.
February 2002 – 333 recordings cataloged.
August 2002 – Began proofreading the finding aid.
August 2004 – Finished compiling the finding aid.
January 2006 – Audio and image files linked to the finding aid.

Imaging

August 2001 – Planned imaging workflow for all three collections; Boulton test scans sent to Digital Imaging Group; images inventoried and organized; selected items for delivery through Page Delivery Service
February 2002 – DIG completed processing of Boulton images; test scans submitted for Ellington record labels.
August 2002 – Completed imaging of Ellington materials; *Image Digitization Questionnaire* submitted for Rubin.
February 2003 – Began imaging of Rubin materials.

Audio Reformatting

January 2000 – Discussed structure of an application-specific Audio Interchange File Format (AIFF) sound file.

April 2000 – Installed *Audio Precision* Sys 2522 Dual Domain Audio Analyzer in studio; wrote small Mac PERL script to manage production of audio file; several sound files run through entire production cycle to gain sample data.

Boulton

July 2000 – Created high density archival and CDR copies for 15 Boulton reel-to-reel tapes.
June 2001 – Boulton high-density and CDR copies completed; began production of Real Audio files.
February 2002 – 70 Real Audio files deposited on the UIS *Visual SourceSafe* server; all named in the NRS and linked to the finding aid.
August 2002 – 70 Real Audio files *moved* from UIS server to the OIS Real Audio server; NRS names updated to reflect the change.
August 2004 – Began deposit of audio files.
February 2005 – Audio deposit complete; files linked and available through OASIS; preservation audio files stored in DRS.

Ellington

June 2001 – Began reformatting of selected recordings.

August 2002 – Completed Real Audio files; considered security and reporting options for ASCAP licensing requirements.
February 2003 –Submitted test file for assessing security of Streaming Delivery Service (SDS).
August 2004 – Began audio deposit.
February 2005 – Completed audio deposit; files linked and available through OASIS; preservation audio files stored in DRS.

Rubin

August 2001 – Created a small number of high-density and CDR copies; altered audio reformatting workflow to give project cataloger more immediate access to Real Audio files.
August 2002 – Began reformatting for streaming audio copies.
August 2004 – Began audio deposit; continued reformatting of materials.
January 2006 – Completed audio deposit; files linked and available through OASIS; preservation audio files stored in DRS.

Software tools

February 2003 – Explored *xorro* xml editor as application for audio object and process history tools.
June 2004 – Audio object and process history tools tested and functional.

Selection/Permissions

Providing digital access is necessarily a selective process. For one, it is rarely possible in financial terms to reformat and deliver every item in a given collection. The common misconception that everything will eventually be digitized certainly flies in the face of this reality. Other reasons can be ethical, when issues of copyright are ambiguous, or legal, when dealing with commercially viable recordings. They might also be technical, as when materials have deteriorated to such a degree that they cannot be restored.

Our basic approach was to reformat as much of these three collections as we could, realizing that not every sound file and image would be deliverable. Our users determined the priorities for which multimedia materials were selected for inclusion in the finding aids. Decisions were made based upon the priorities of Harvard music faculty and students as well as the larger research community, and on the strengths of our collection. Our criteria for content were uniqueness and diversity of interest. In a practical sense, we also had to preserve some deteriorating open reel tapes first, and then make the access decisions.

In seeking permissions for both the digitization and networked access of these rare materials, we encountered a variety of intellectual property scenarios. Each collection had its own challenges and concerns, and often the copyright work was very time-consuming. On a case-by-case basis we consulted with the University's Office of the General Council to determine the rights and fair use issues regarding our collections.

The University owned the rights to recordings in the Boulton collection and there were no real issues in making them publicly available. However, the Ellington 78s and LPs were commercial recordings and most works recorded by the great artists in the Rubin collection were commercially viable. In the case of Rubin, the recordings were over 40 years old and from a time and place when signed permission forms were not commonly in use. While we did seek permission from performers or their descendants, we could never be completely certain what kind of permissions Rubin himself had been granted. With no attempt to provide access to the whole collection, we focused our efforts on three South Indian performers. Initially seeking oral

agreements, we ultimately carried out all legal work by letter and were given permission to stream versions of their performances.

For the Dodge collection we selected a small number of pre-1972 recordings and entered into negotiations with ASCAP/BMI for the rights to digitize and stream those. No formal licensing was required for BMI and all work with ASCAP was complete by 2001. Because of our agreements with ASCAP, we had to make special efforts to guarantee the online security of the Ellington recordings. In August of 2002, our progress on the Ellington finding aid was interrupted as we waited for the security and statistical reporting capability of the Real Audio server. OIS eventually created the necessary servlet and in early 2003 we submitted a test file for assessing the security of its Streaming Delivery Service (SDS). Once this was confirmed, the Ellington files were deposited and linked.

After working out copyright issues for the collections in LDI 2, some questions still remain. What type of standard agreements should we choose to have in place? Is there even a point to having them? And regarding rights, should we work toward an expansion of fair use instead of attempting to gain rights to all the material in our growing collections? And if we choose the fair use, what collections and which tracks do we take the calculated risk of putting up?

Finding Aids/Supporting Materials/Imaging

One of the primary goals of LDI 2 was to create unique multimedia finding aids that integrate sound, image and text into a single electronic research tool. By doing this we hoped to provide scholars with the means to study materials that might otherwise be too fragile or remote to access. Since this project began in 1999, such rich multimedia electronic resources have in fact shown themselves to be extremely valuable to scholarly research.

We chose to create these resources using Encoded Archival Description (EAD), at the time already a standard for encoding electronic archival finding aids. EAD was a DTD (Document Type Definition) that had originated in the archival community and was primarily designed to describe text-based materials. However, we believed it to be more flexible than standard cataloging, and hoped to use EAD not only to integrate image objects but also digital versions of our sound recordings. At the time, EAD had not yet been used for audio collections and it was a stretch to imagine it could be used to describe time-based media.

Traditional archival finding aids deal with physical, printed collections that are often arranged in predictable ways. However, in approaching our three audio collections, we were faced with more challenging decisions about how to organize the materials. Series heading attributes in paper-based finding aids generally describe the physical objects, but with these collections we had to be more creative. Thus, the Boulton collection was organized by the geographic location of her recordings, Dodge by recorded format and Rubin by field trip and in chronological order. The Rubin finding aid provided the greatest challenge because there was no obvious way to break down the more than 1000 pages of material. The OASIS group proved very helpful in this process.

In the world of printed matter, published works are generally catalogued in OPACs, while unpublished works often appear in archival finding aids. However, such divisions are not always so clear in field audio collections where unique, unpublished recorded material often sits side by side with commercial or at least commercially viable recordings. The first impulse

might be to keep the collection intact, but there can be problems with this approach. With the Dodge collection, which is made up of mostly commercial discs, we should have done more item-level work to cover the track listings for each of these recordings. One could argue that this material might have been better served by traditional cataloging.

Once inventories were taken for each collection, the finding aids were assembled and we began the mark-up process. We chose to work on the Boulton finding aid first because the collection was smaller and more manageable, and because its organization by tape was simpler than the others. This proved to be a good decision, and the initial phase of the project proceeded smoothly while our newly hired project assistant learned the SGML necessary to work with EAD. With Boulton, we got a good idea of the basic processes and later we were easily able to apply what we learned to the other two collections. This became the first finding aid to be completed and loaded into Harvard's OASIS system.

Much early work on the Boulton finding aid was done before the advent of current technologies and systems. For instance, the finding aid was initially marked up before our computers had browsers that could correctly display SGML documents. Not surprisingly, this led to a lot of extra work and a number of tagging and formatting errors. We also scanned Laura Boulton's logbooks before the Page Delivery Service (page-turning) system was in place, so we were unable to see how it would eventually work. Once PDS was up and running, we had to revise our early descriptions of the logbooks to make them work with the new system.[2]

Working with supporting documentation from the Boulton collection brought us into contact with Harvard's Digital Imaging Group (DIG). In the summer of 2001, a number of meetings were held between project staff, HCL DIG and OIS to plan the imaging workflow for all three collections. It was determined that each image would exist as a graphic interchange format (GIF) and a tagged image file format (TIFF) file. At that point the repository was not available for storing our images, so we agreed that no scans would be sent to HCL DIG until the DRS was ready to accept the files.

[2] Adams, Sarah. "EAD for International Music: Digital Access to Archival Collections." Paper presented at IAML annual meeting 2003, Tallinn, Estonia, pg. 4.

Work with Boulton materials was straightforward, though we did have to bring some earlier scans up to Digital Fine Art Printing (DFAP) standards. These had been done prior to the existence of the DRS or HCLDIG and their metadata had to be coerced into an XML loader document. A one-off conversion script for loading these scans was created to address the issue.

Later, test scans were sent to Imaging Services and it was agreed that the master TIFF files were of high enough quality for on-screen display. Images from the Boulton collection were then inventoried and organized for submission to DIG. We also met to decide which parts of the collection would be delivered through the new Page Delivery Service and how to structure the metadata for different types of documents. Selected Boulton materials were then sent to DIG so that curators could test the new PDS interface.

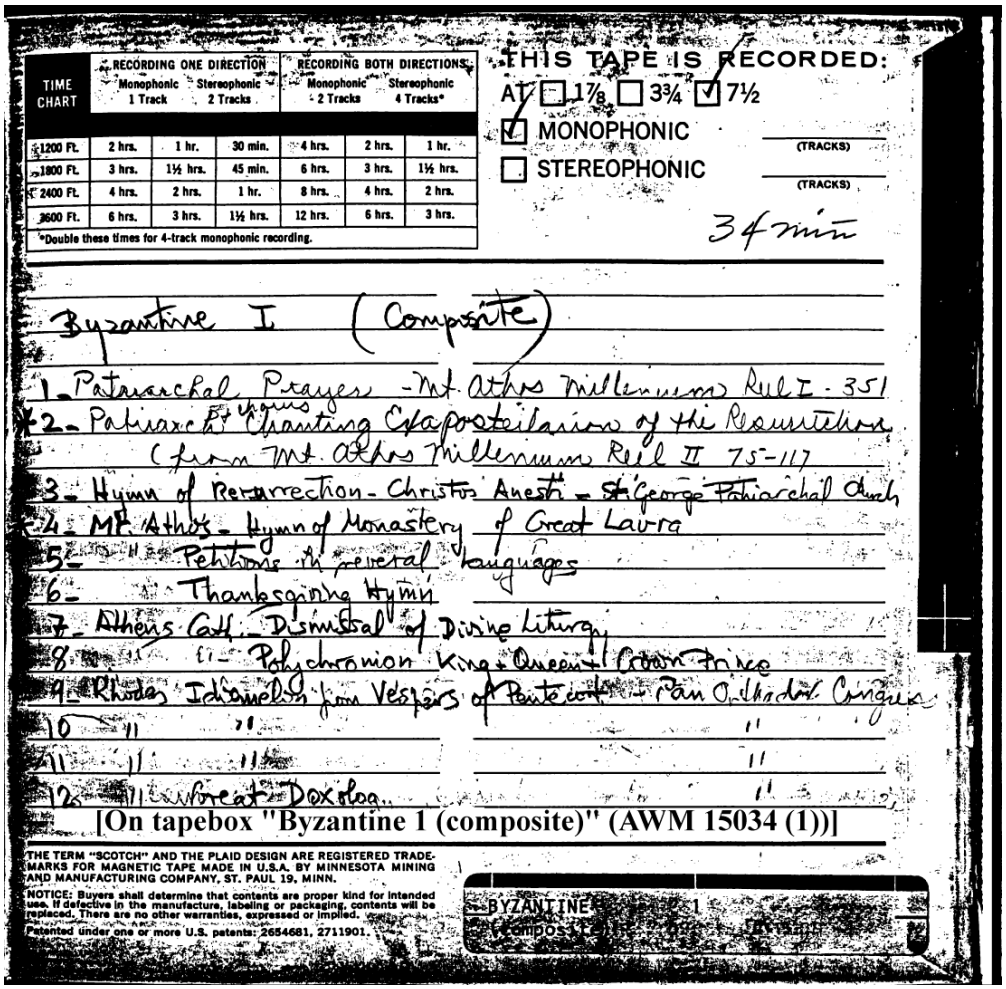

Figure 1. Scanned image of notation on a Boulton tape box

Laura Boulton's collection contained notes, logbooks, tape boxes and other miscellaneous documents (*figure 1* is an image of one of her tape boxes). Scanning for these items proceeded without any real difficulties and, once it was complete, we moved on to the Dodge collection. Here, materials to be scanned consisted mainly of 78-rpm record labels and a few other supporting documents. As with Boulton, we began with test scans before sending any 78's over to Imaging Services. Ultimately, work on the Dodge materials flowed even more smoothly due to the quick turn around time for processing these record labels.

The first metadata model we used for images in the Boulton collection was MOA II, a standard from an earlier Digital Library Federation project called the Making of America II. This XML DTD defined the elements and metadata (descriptive, administrative and structural) encoding levels for each digital object. The project assistant described images with metadata pages that were then supplied as directions to the photographers at Imaging Services. It was a very time-consuming process because each photo needed its own page of metadata instructions.

Images were then sent to DIG for processing where the MOAII metadata was created based on our specifications. Using Uniform Resource Names (URNs), each object then had to be named for its persistent linking from the finding aid. The imaging lab also provided the NRS names using the url-3 designation for each MOA2 document. Once the URN was plugged into the coding of the finding aid, the link then called up the digital object. It would later work the same way with our audio files.

In the summer of 2002 we held meetings to discuss transitioning our metadata from MOA2 to the Metadata Encoding and Transmission Standard (METS) schema. We had many basic questions. How would METS work an as organizing structure for our purposes? How would we create these METS documents, particularly DRS and AES (Audio Engineering Society) metadata for sound files? Once plans were made, we posted our ideas about structuring METS data to the METS listserv for discussion and comment. We received no negative feedback.

The relatively small size of the Boulton collection meant that assembling entries and proofreading text in that finding aid was a fairly easy process. But although the Ellington finding aid began as a simple Word document, it did eventually grow to be quite large. It was here that we first encountered

issues relating to the size of finding aids. Because of its even larger size the Rubin finding aid presented the greatest challenges, often pushing the limits of WP8, the version of Microsoft Word we used to create this finding aid.

The Rubin finding aid ultimately even became too large for OIS to manage and at that point it had to be broken down into around 30 separate sections. These sections were then stitched together in the final display. Overall, we spent an enormous amount of time and energy preparing this finding aid, and much more clerical work was required than we had originally thought. It has been suggested that when working with this much material in the future, logical divisions should be made in advance and each section should make up its own finding aid.

The content of the Rubin collection also placed unique demands on the project. We knew the material in this collection would require a specialist with a broad knowledge in South Indian classical music and significant Indic language skills, but we were not prepared for the amount of time and effort it took to find the right person. We did finally hire our project cataloger late 2000 and one of the first things we did was train him in basic cataloging practices and EAD. Predictably, the Rubin materials required considerable original cataloging, and our cataloger routinely spent many hours listening in order to accurately describe the musical content and determine access points. A basic grasp of technology was also necessary for this position as the cataloger frequently worked with our audio engineer to identify content divisions during the indexing process.

Still other reasons contributed to our very slow progress in assembling the Rubin finding aid. These included punctuation problems, the rampant need for normalizing names, and general difficulties in handling the diacritics of several Indic languages. We faced significant authority issues because materials in the Rubin collection demonstrated considerable variation in proper names. Sometimes this was due to differing transliterations of a given Sankrit or Tamil name, but other times it was simply due to inconsistencies in the supporting documentation.

Very few Library of Congress headings could be found for performers in the Rubin collection. We used the authority file version when it existed, but in some cases we chose the Roman script versions preferred by the performers. In each case a studied version had to be arrived upon by comparing Rubin's original version, the authority file, and the thoughts of the project cataloger.

The extensive cross-referencing necessary to deal with all of these variations also created a huge amount of work. Once again, we had completely underestimated the time and effort that the considerable authority work would require.

We also learned that EAD rarely provided for the unique elements found in ethnographic collections. Diacritics and special characters in Indic languages consistently created display problems that more often than not required consultation with OIS for solutions. We ultimately dealt with many problem characters by creating electronic entities to represent them. But again we were surprised at how much attention had to be given to correcting display issues. One final issue surfaced in early 2006, just after the project associate finished linking audio files to the Rubin finding aid. In this case a number of Unicode characters failed to display properly, and again OIS provided the additional processing necessary to correct the problem. Looking ahead, the challenges we faced with diacritics gave us a good indication of what to expect when working with anything other than a standard alphabet.

Ironically, we also saw substantial variation between nicknames and given names of jazz players featured the Dodge collection. Although it was not the ideal solution, we ultimately standardized the forms of names throughout the document. We hope to improve access in the future by providing a key to variants and/or links to authority records.

Proofreading documents as elaborate as the Rubin and Dodge finding aids also took much more time than we had anticipated. Because of its considerable length and other idiosyncrasies, Rubin was a particular nightmare. When the proofing finally was complete, we then assembled Rubin's many separate parts into one master document and began to integrate multimedia links and diacritic entities. Fortunately, all the normalizing work we had done on this finding aid had served to make the text more consistent for hyperlinks. Aside from a couple of minor link syntax issues, the linking process went smoothly and we continued to add fresh links for newly digitized audio files throughout the remainder of the project. In fact, the successful establishment of links from the finding aid to its digital objects and sound files represents one of the quiet successes of this project.

Rubin turned out to be a massive undertaking with so much detail and so many challenges to everyone that we did not complete the finding aid,

cataloging, reformatting and linking until five years after the project had begun. At one point, to help the process along, we even adjusted the audio reformatting workflow to give the project cataloger more immediate access to the Real-Audio files.

The volume and complexity of supporting materials in the Rubin collection was more substantial than either Boulton or Dodge. A wide range of reel-to-reel tape boxes, programs, loose notes and notebooks ultimately produced around 1000 images. The project assistant had to sift through a bewildering number of document descriptions in preparing these materials. Below is scanned image of a notebook from the Rubin collection (*figure 2).*

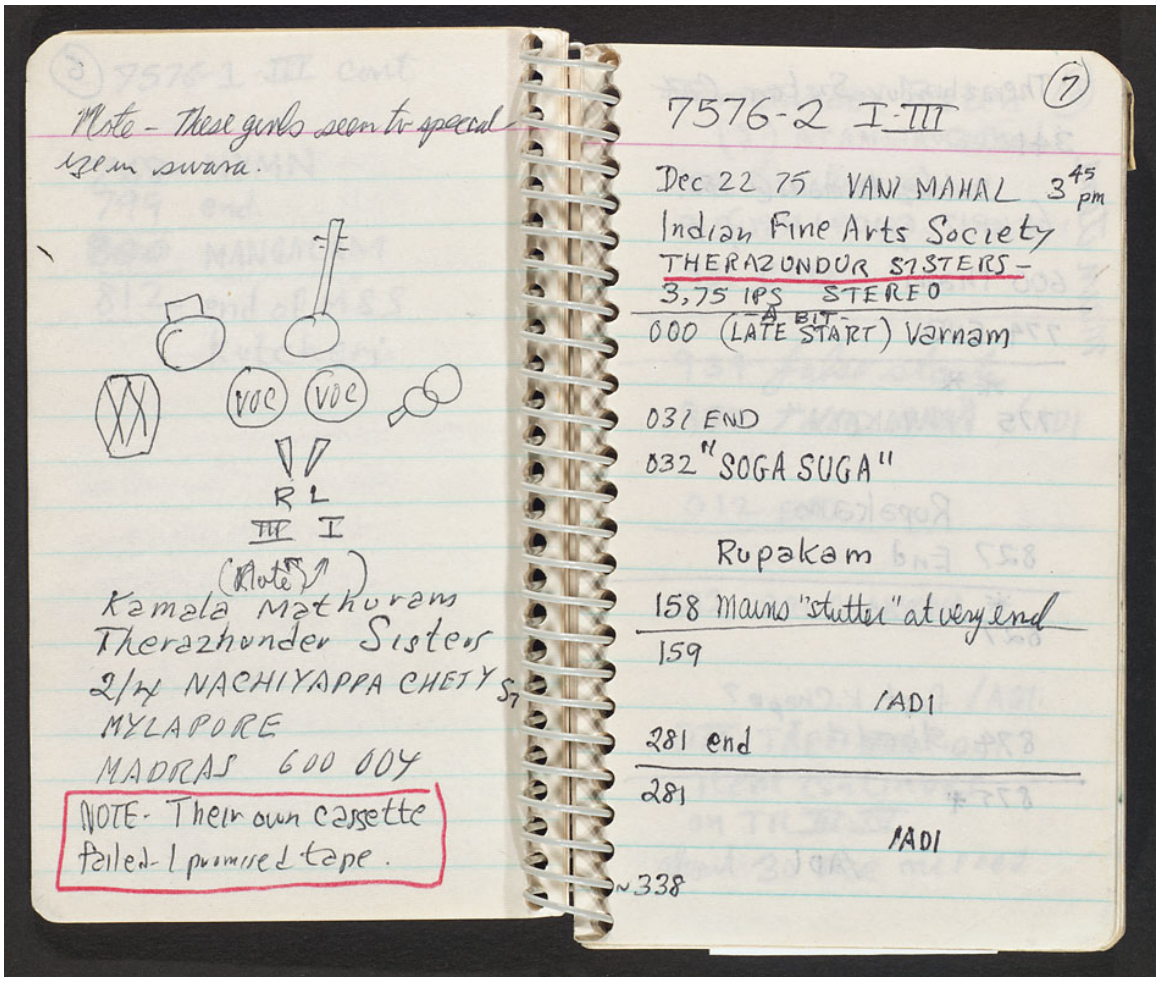

Figure 2. Scanned image from a Rubin notebook

We learned many lessons from building finding aids for these three very different collections. Perhaps the most obvious is that we would not again choose to create a finding aid anywhere near the size of Rubin without breaking it down into smaller sections first. We also need to be ready at all times to respond changes in format. XML is now the preferred language of EAD, but when LDI 2 began the standard markup language of the day was SGML. All three of our finding aids were originally created in SGML and later had to be converted to XML using Word Perfect 12. Also, the preparation of these finding aids began under the first version of EAD. With the advent of the 2002 update, some of the initial attributes had been discontinued, and the finding aids manager had to make the earlier work fit into the new version.

A highlight of this project was the successful integration of differing types of digital objects into a single finding aid. However, in the future we hope to go even further by creating a single digital object that contains within itself multiple manifestations (audio, text, image, etc.). Such an approach would give us an even greater ability to show relationships among all the parts of our collections. [3]

[3] Adams, Sarah. "EAD for International Music: Digital Access to Archival Collections." Paper presented at IAML annual meeting 2003, Tallinn, Estonia, pg. 2.

Audio Processes

a. Preservation/Reformatting

When the *Music from the Archive* project began Harvard had no audio preservation program, and none of the University agencies or services now associated with audio delivery existed either. And not only were we starting from scratch at Harvard, but we were also moving forward in a process for which there was not yet even a national standard.

In late 2000, our project manager attended the Audio Engineering Society conference and presented a paper that detailed all the audio metadata decisions we had made up to that point in order to initiate work on LDI 2. He then took a lead role in the in the larger discussion, eventually contributing two influential AES metadata standards.

In the very beginning we spent a lot of time just planning infrastructure and working out details. Some of the basic decisions we had to make before any reformatting could be done included choosing the type of audio delivery format, the location and ownership of the server(s), and how we would use the repository. Having previously used analog tape or DAT, we also had to imagine a reliable delivery system at the heart of this process. And, as we made our way without an established framework, it was important that we maintained the insight to know where our strengths and weaknesses were.

Early discussions centered around how we might create an application-specific AIFF chunk, what its structure would be and what possible metadata fields might be necessary. This chunk would store all the metadata and allow anyone in the future to recreate the original sonic processing files, possibly substituting newer processing algorithms on a selective basis. After we were satisfied with the feasibility of this approach, we started work on a small Perl application that would both manage the process of building and embedding the chunk and recover the metadata from the sound files.

Next we outlined the reformatting production cycle. Each audio file would exist as a high-density Sonic audio interchange file format (AIFF) file, a

mid-density vanilla AIFF (16bit/44.1khz) and as a real audio file (G2 T1, 56k, 28k). Then several sound files were created and run through the entire production cycle and no major problems were encountered. While we were testing the metadata application, files and data were backed up using traditional backup utilities. Once we were satisfied that the application was producing good data, additional backup was discontinued.

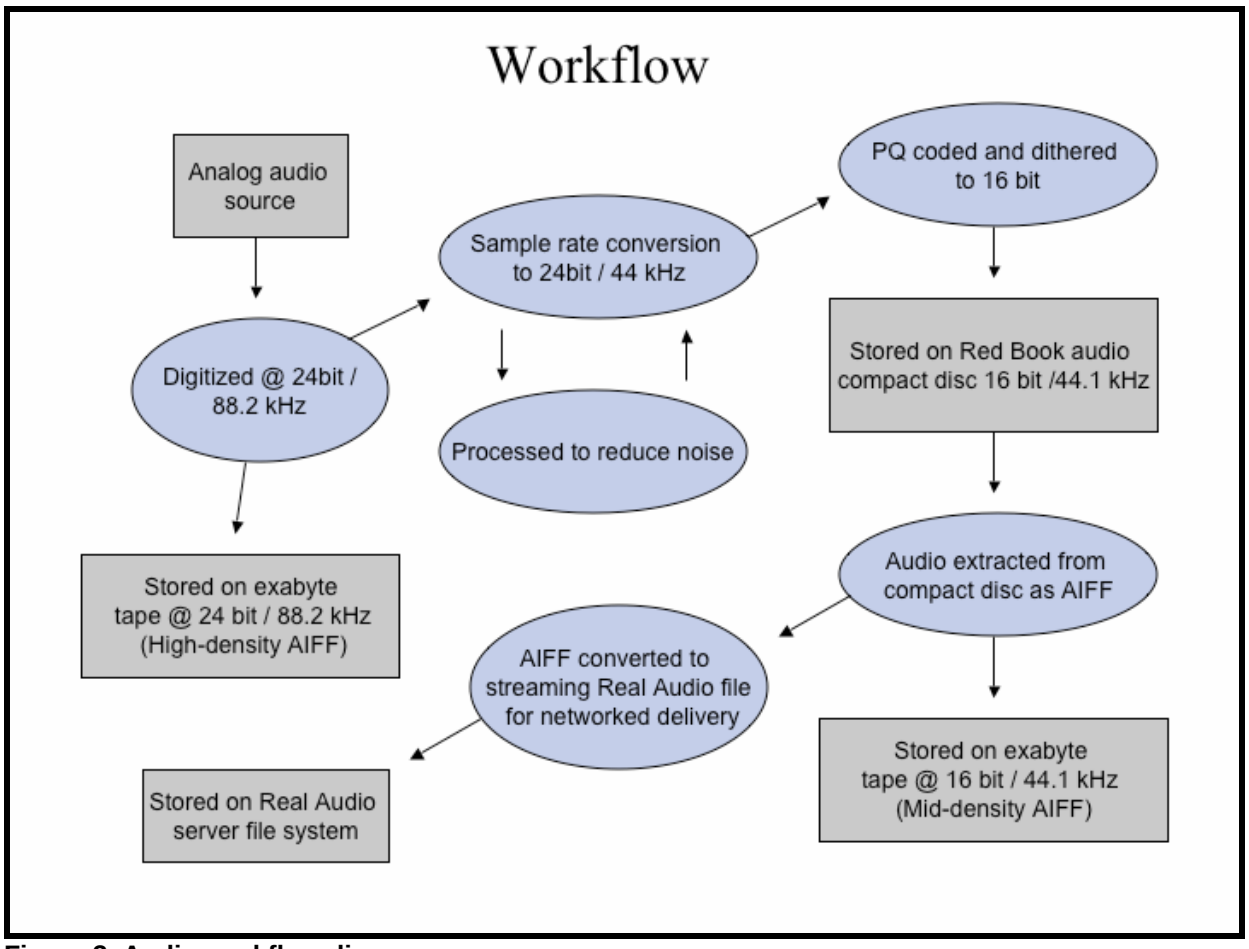

Figure 3. Audio workflow diagram

This trial process evolved into the full-fledged production cycle we used throughout LDI 2. The high density AIFF *archival* file was created first, from which we derived a CDR working copy. From that copy we derived the mid-density AIFF, which in turn was used to derive the final Real Audio file. An outline of our audio reformatting process appears in the workflow diagram above *(Figure 3)*.

Many of the open–reel tapes in the Rubin collection had degraded much more than we had originally thought. For this reason the reformatting

process moved very slowly and handling time doubled from our original projection of three hours of labor per one hour of resulting audio. To restore these recordings, we needed to capture a large number of edits and then reassemble them as one continuous file. Fortunately, given this extra attention, material on the original tapes was recovered.

Once the Real Audio file was derived we had to store it somewhere. In considering our storage options for these sound files we came to the first major lesson of the project. The current file storage infrastructure of the DRS was woefully inadequate for accommodating the large size of audio files. Next we had to come up with a reliable way of retrieving the sound file. Thus, another important goal of LDI 2 was to build a delivery system for the DRS.

The basic unit of that delivery system was going to be an audio submission information package (SIP), a container that defined the way audio and metadata would be combined for deposit into the DRS. We had many questions about how to structure these SIPs and how the DRS would be set up for ingesting them. The first task at hand was to specify the SIP's core metadata attributes. However, knowing how to define the elements was not always clear, and building the tools to create/read the metadata took several years. In retrospect, developing a SIP for audio was anything but a straightforward process, though our work here was ultimately one of the project's major successes.

Absolutely essential to this project and our goals for a continuing program was our plan to develop an automated method for collecting metadata. In the absence of such a tool, two new software programs for collecting metadata were developed during the course of LDI 2. One gathered digital provenance metadata and the other technical metadata for audio file tapes. Both are available to anyone by request. When we first started work on the *Music from the Archive* grant, the metadata model we used was MOA2. Later, after much planning and discussion we made a transition to METS an as organizing schema. In the process we had to pay special attention to how DRS and AES metadata would be created in this new format.

We also aimed to capture a fair amount of structural metadata for relating sound files to one another. For this we turned to the Audio Engineering Society's AES-31-3, a standard for the networking and transport of audio file formats. But as this was all new work, the road to getting it right was a

protracted one at best. Progress depended on the infrastructure and execution of several Harvard agencies, many of which were also new to this work.

One of the most basic challenges in dealing with field recordings can be determining the beginning and ending points of performances on tape.

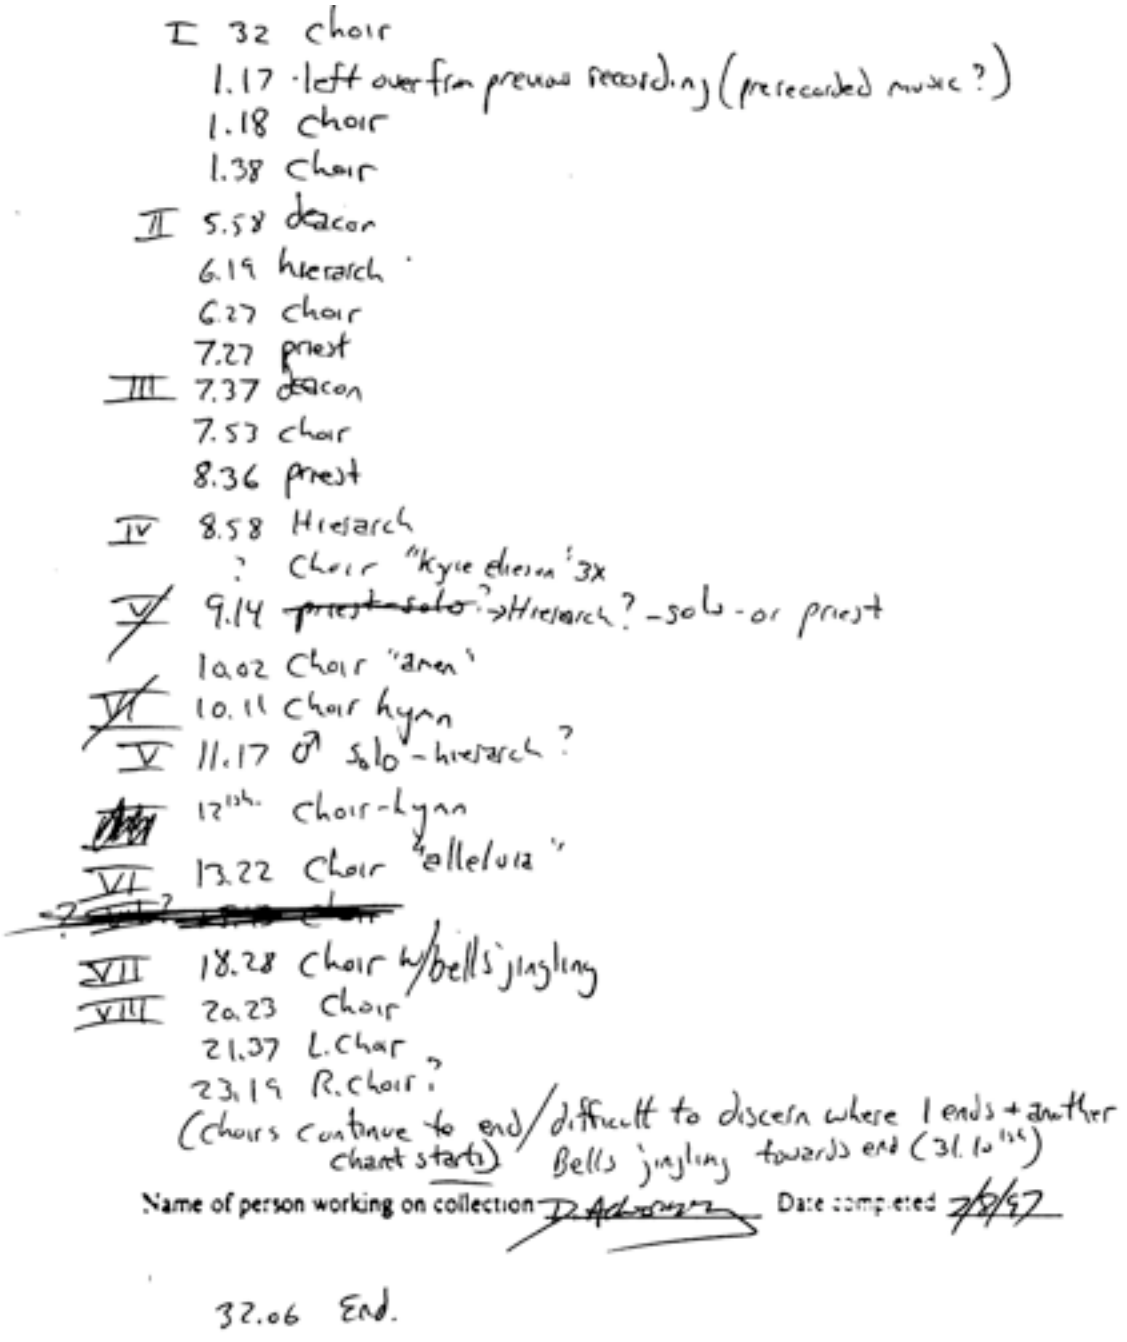

Figure 4. Boulton audio worksheet

It’s not unusual to see timings for recordings in these collections that are significantly off, and we certainly found this to be the case with many Rubin recordings. Doing content division with such inaccurate information was often very frustrating. Above is an early Boulton audio worksheet that illustrates the point *(Figure 4)*.

Another related issue we encountered with the Rubin collection was that of individual discrete sound files**.** Original audiotape selections normally contain multiple performances. When that tape is converted into digital sound files, unexpected overlaps and breaks sometimes occur, leaving no correlation between sound file and performance. The large number of edits needed to reassemble many of the Rubin recordings only contributed to the problem. Thus we relied on structural metadata to maintaining accurate mapping and for showing the relationship between all these parts. This often slowed the Rubin tape reformatting process dramatically, but fortunately we would later solve the issue in *Sound Directions*, a NEH joint digital preservation project between Harvard and Indiana University.

Two other issues affected reformatting of the Rubin recordings. The first was ultimately due to the sheer amount of materials (nearly 2000 hours of never cataloged performances) as well as the time-intensive nature of working with them. It involved shifting the audio reformatting focus to give the cataloger access to the recorded material first. This was done so that the Rubin finding aid, whose myriad issues have already been discussed, could finally be finished.

To make this work, we began to produce the high-density and mid-density AIFFs concurrently, followed by a Real Audio file for cataloging purposes. Since the metadata could not be finalized until the CD copy was produced, the bit streams had to be stored on computer tape until we could return to them for any necessary editing and cleanup. At that point we finished the metadata and wrote the final archive master. This deviation from our original workflow also affected the deposit of Rubin audio files into the DRS. Because of the time it takes to restore the audio files from the tape archives, the process of retrospectively depositing was more time consuming. By looking at these scenarios alone, one can see why our timeline was at times necessarily eccentric.

The second had to do with providing reliable supporting documentation (metadata) for the reformatting process. In certain ways we were lacking in strict guidelines for how to document materials, and some of the resulting information was problematic. As the implications can be far reaching, making sure documentation is in order should be an essential part of any reformatting project.

A key component of the quality assurance process in our studio is an *Audio Precision* Sys 2522 Dual Domain Audio Analyzer that was purchased at the beginning of the project. This system is used for proof of performance testing. Our *Prism* AD2 Converter, a cornerstone of our studio, can capture original audio from 78s or open-reel tapes at a very high rate (88.2 kHz sample rate) and in superb detail (24bit digitization). At the outset of LDI 2 the audio preservation studio also physically underwent a major acoustical upgrade. As a result the listening environment became much more neutral, making for better decisions in the preservation and production process.

Besides successfully creating a body of digitized audio recordings, another direct result of the project has been an increase in the speed of our connection. We now have a 1-gigabit Ethernet for the throughput of data on our network, miles ahead of where we were when LDI 2 started. We've also made significant strides with storage capacity and the delivery of streaming audio, something that began with this project. Finally, the technical metadata collection program was a very good investment because it is now used heavily in most all of our audio preservation work.

b. Software Tools/Deposit/Delivery

Imagining and developing a system for delivering digitized audio, where there was none in place, meant working on several interlocking issues simultaneously. We first needed an appropriate sound file format, then a reliable system for delivering these sound files, and finally, we needed automated tools to help gather metadata and facilitate the deposit of these files into our repository.

Real Audio was chosen early on to be the most suitable format for our deliverable sound files, and in the summer of 2001 we began discussing possible delivery systems for dynamic/complex multimedia objects. At that

point OIS was not able to devote resources to the development of such a system. As an alternative, we then considered hiring a contract programmer to develop a web interface for our collections. We went so far as to create a mockup of a multimedia style sheet to evaluate the feasibility of that approach.

At the same time we had also begun development of the two software tools that would be crucial to the work of LDI 2. Once completed, the core audio (later renamed audio object) tool would facilitate the deposit of reformatted audio files into the DRS, and the APXE (Audio Processing XML Editor) would capture metadata on the work history of each sound file. Using the XML editing application *xorro* as a starting point for these tools, we completed the audio object tool in 2003. However, the process history tool took much longer and we ultimately hired an outside programmer to see it through development and testing. Concurrently, OIS was also working on DMART (DRS METS Archive Tool) an automated tool for assembling METS and DRS deposit documents.

Our initial plan was to begin the deposit of Boulton and Rubin audio files before finishing the process history tool. However, we ultimately decided to wait on the creation of most deliverable audio files until the process history and DRS load software tools had been completed and tested. By early 2003 DRS deposits for all three collections had been held up for some time as we awaited these tools. Target dates for the deposit of Real-Audio files repeatedly came and went due to both collection-specific factors and purely technical challenges. In order to make it work with the process history tool, we also had to make changes to our original audio object tool schema.

By the summer of 2004 audio deposit was finally underway for the Boulton collection. Once again, though, we hit a snag as OIS discovered a problem and all deliverable files had to be re-deposited. Fortunately, OIS updated the NRS by pointing the currently assigned names to the new files, thereby saving library staff from having to update Boulton finding aid links. By early 2005, the deposit of Boulton and Ellington files was complete and those recordings could finally be accessed through finding aids in OASIS. The library then hired a second audio engineer who was trained in the deposit of the remaining Rubin sound files.

We moved through a number of scenarios in concert with multiple Harvard agencies before arriving at a viable solution for delivering Real Audio files.

Initial discussions in 2000 went smoothly as we determined the metadata specification required for Real Audio. But when faced with the question of who would serve these files, the currently available choices were limited. Our first plan was to have FAS serve LDI 2's real audio files. However, the huge size of these files would have demanded the purchase of additional streams and storage space that made this approach impossible. Next we discussed the question with University Information Systems as it had recently developed a streaming audio service. Initial cost estimates were prohibitive and we had doubts about whether UIS could offer an affordable cost structure for our needs. A final option was to have OIS purchase a Real Audio server for the delivery of these files. Ultimately, this is what happened.

In the short term we had decided to place our early real audio files on the UIS Visual SourceSafe (VSS) server, but a number of issues put this system at odds with our program. UIS had a script in place for auto-deleting files that were not touched for three months. Although it was supposed to have been modified to keep our files intact, we ran into problems with this script. The VSS system also lacked any mechanism for owners to delete files and prune the collection, and most importantly it had no reporting functions. The latter issue made it impossible for us to keep our ASCAP and BMI licensing agreements for the Ellington recordings.

During this time OIS did purchase its own Real Audio server. And in early 2003, the Streaming Delivery Service (SDS) was up and running for the express purpose of delivering real audio files stored in the DRS. In light of theses developments, seventy real audio files were moved from the UIS VSS to the new server at OIS in August of 2002. NRS names were then updated to reflect the new location. We had to wait a bit longer for the security and statistical reporting capability necessary for the Ellington recordings. But once that was done, the sound files were linked and the Ellington finding aid completed.

Budget

Final budget figures for LDI 2:

Salary and Fringe Benefits = $172,369.87
Imaging services: scanning, structural and administrative metadata and DRS deposit (HCL Digital Imaging Group) = $11,730.80
Outsourced Computer Programming (including design, development and testing of Metadata Collection Tool for Process History) = $19,460
Audio Consulting = $160

Total Expenditures: $203,720.67 (Project cap: $204,000)

Personnel:

As one would expect, labor costs accounted for the largest share of the project's overall expense. The four core positions of LDI 2 were:

- audio engineer

  Project Manager: oversaw the project; wrote two influential standards for collecting and managing audio metadata; worked closely with OIS to develop new technology for producing multimedia documents; designed and created software tools to further the work of this and future projects; converted original audio sources into digital format; worked with catalogers to identify content divisions during indexing process; worked with curators during the selection process.

- finding aid operations manager

  Finding Aid manager (Project Associate): supervised the production of all electronic finding aids; trained staff in EAD; consulted on design issues; oversaw markup and the linking of supporting materials; ensured adherence to standards; adapted earlier work from first version of EAD to the 2002 update.

- digital project assistant

  Project Assistant: prepared materials for scanning and deposit into the DRS; assisted in the creation of Real Audio files for all three collections; marked up lengthy documents with SGML tags; used SGML to link sound files and other supporting documentation to the finding aid; provided assistance with naming conventions, authority file issues and diacritics.

- subject specialist

  Subject specialist (Project Cataloger): provided subject expertise in deciphering and describing the content of Rubin field recordings; catalogued materials after being trained in EAD and cataloging methods.

  Note: We consulted with one other specialist for a small number of Boulton recordings, and also drew on the knowledge of departmental staff members for some materials.

Imaging:

In order to give an idea of imaging costs for the variety of supporting materials scanned (logbooks, program books, tape boxes and inserts), below is a sample billing record from the HCL Digital Imaging Group *(Figure 5)*.

Project Summary (details for each category on following pages)

Category 1: spiral notebooks	**product**	**type**	**$ per unit**	**total #**	**total $**	**total MB**
	page images	archival_master	$7.00	1,266	$8,862.00	18,990
		delivery_1	n/a	1,266		247
	metadata	for pages				
		for object	$5.00	32	$160.00	n/a
SUBTOTALS				2,564	$9,022.00	19,237
Category 2: program books	**product**	**type**	**$ per unit**	**total #**	**total $**	**total MB**
	page images	archival_master	$7.00	26	$182.00	133
		delivery_1	n/a	26		6
	metadata	for pages				
		for object	$5.00	2	$10.00	n/a
SUBTOTALS				54	$192.00	139
Category 3: tape box inserts (> 12")	**product**	**type**	**$ per unit**	**total #**	**total $**	**total MB**
	page images	archival_master	$15.00	26	$390.00	962
		delivery_1	n/a	26		11
	metadata	for pages				
		for object	$5.00	26	$130.00	n/a
SUBTOTALS				78	$520.00	973
Category 4: tape boxes and inserts (<= 12")	**product**	**type**	**$ per unit**	**total #**	**total $**	**total MB**
	page images	archival_master	$7.00	200	$1,400.00	2,400
		delivery_1	n/a	200		44
	metadata	for pages				
		for object	$5.00	62	$310.00	n/a
SUBTOTALS				462	$1,710.00	2,444
GRAND TOTALS				3,158	$ 11,444.00	22,794
DRS Annual Storage Cost	**$111.30**			(current rate = $5.00/GB ... 1,024MB = 1GB)		

Figure 5. HCL DIG summary billing record

Storage:

Similarly, the sample costs below indicate the significant costs incurred in storing copies of sound files on exabyte tape (archival and mid-density) and in Real Audio format.

Archival copies:

210hrs x 30M/min = 370G
370G x $18G = $6,660/year to store offline

Mid-Density copies:

140hrs x 5M/min = 41G
41G x $18/G = $738/year to store offline

Real Audio copies:

140hrs x 34M/.5hr = 9.5G
9.5G x $520/G = $4,940/year to store offline

Total for All Three Copies:
$6,660 archival
$738 mid-density
$4,940 delivery

$12,278/year

Total for Just Archival and Delivery Copies:
$6,660 archival
$4,940 delivery

$11,540/year

*Purchased by HCL for the project, a key component of our studio is an *Audio Precision* dual domain analyzer. The final budget includes $160 in calibration and consulting fees related to the set up of this machine.

Conclusion/Project Benefits

One of the most significant legacies of this project is that it gave Harvard an audio preservation program. It's also worth noting that the tools created in LDI 2 to support that program have already been put to good use in several subsequent projects, notably *Sound Directions*, an NEH joint digital preservation project between Harvard and Indiana University. We also have a small group of staff members taking what was learned in *Music from the Archive* and applying it to more recent projects in the LDI program. In these ways we seem to be well on the road to institutionalizing some of the basic practices developed in LDI 2.

And what of EAD for audio? Our work in LDI 2 has shown that, although it was not intended for use with sound recordings, EAD can be a viable DTD for the expression of audio collections. General functions and tagging relationships basically fall into place, although one can always hope for further refinements in the future. So far we have already seen a new element for *material-specific details* appear in the 2002 EAD update. This attribute nicely accommodates playing time for sound recordings.

And what of our goals for preservation and expansion of access to unique and rare sound recordings? Over the course of LDI 2 we rescued a large body of significant recordings from potential oblivion, and produced nearly 9 gigabytes of digitized Real Audio sound files. These recordings are now delivered through stable finding aids that can easily be accessed by interested researchers and members of the Harvard community. Are they *useful*? We do know that the Boulton finding aid has been a very popular resource for our patrons.

Thus far, have we created *durable* digital audio formats, metadata systems and storage facilities? Time will tell, but we do know that AIFF and Broadcast Wave have shown themselves to be preferred and stable audio formats, EAD remains the XML standard for encoding finding aids, and the Digital Repository Service is now well established for storing audio and delivering it through the Streaming Delivery Service. And have we achieved managed storage, where data management, archival storage, administration and preservation planning are all part of the same equation? The DRS now not only offers electronic storage for digital objects, but it also provides

management of administrative and structural metadata, as well as preservation policies to ensure usability over time.[4]

And who is needed to do this work? Based on our experience, the following core positions are essential for any program aimed at comprehensively digitizing audio collections:

- audio engineer
- finding aid operation manager
- digital projects assistant
- subject specialist

In this report one can find frequent references to phrases like “more difficult than we expected” and “took much longer than we thought.” The fact that we encountered continual setbacks is not surprising when you consider we were creating a program from the ground up. The complexities of tackling several related objectives simultaneously might also help to account for our sometimes eccentric timeline. However, in the end we accomplished a substantial amount of original work, and pointed the way to several important standards and best practices.

[4] Overview: Digital Repository Service (DRS), 25 May 2007. President and Fellows of Harvard College. 1 Oct. 2007 http://hul.harvard.edu/ois/systems/drs/

Appendix (Inventory)

A final audio count:

Nearly 9 gigabytes of real audio files

Sound files created and deposited in the DRS totaling 503 archival masters, 999 delivery files and 89 production masters

An overview of all deliverables:

Images: Each item exists as a GIF and a TIFF file

Boulton Collection

1475 GIF files approximately 53KB each. Includes notebooks, transfer worksheets and miscellaneous documents. All required page-turning capability.

1475 TIFF files approximately 8,150 KB each. Includes notebooks, transfer worksheets and miscellaneous documents. All required page-turning capability.

Ellington Collection

3 78RPM disc center labels (color) – one image each, for side A & B
3 Record jacket liner notes (color)
1 Press release (2 pages, color)
1 half-tone photo with caption
1 map of Harlem night spots (color)
1 color pencil notes inside lid of reel-to-reel tape box

Rubin Collection

600 pages of bound concert program books (color). All required page-turning capability.

29 spiral-notebook logs each containing approximately 35 “openings” (color) totaling 1015 images. All required page-turning capability.
40 reel-to-reel tape box exteriors (color)

Audio: Each audio file exists as a high-density Sonic AIFF, mid-density vanilla AIFF (16bit/44.1khz) and as a real audio file (G2 T1, 56k, 28k).

Boulton Collection

55 hours selected by curator representing 10 to 15 separate services

Ellington Collection

25 78 RPM disc recordings approximately 10 minutes each totaling 4 hours and 10 minutes

Rubin Collection

140 hours representing 70 to 100 separate performances