

“They’re all going out to something weird”: Workflow, Legacy and Metadata in the Music Production Process

Glenn McGarry, Peter Tolmie, Steve Benford, Chris Greenhalgh, Alan Chamberlain

University of Nottingham

Nottingham, UK

{author firstname.lastname}@nottingham.ac.uk

ABSTRACT

In this paper we use results from two ethnographic studies of the music production process to examine some key issues regarding how work is currently accomplished in studio production environments. These issues relate in particular to workflows and how metadata is adapted to the specific needs of specific parts of the process. We find that there can be significant tensions between how reasoning is applied to metadata at different stages of production and that this can lead to overheads where metadata has to be either changed or created anew to make the process work. On the basis of these findings we articulate some of the potential solutions we are now examining. These centre in particular upon the notions of Digital/Dynamic Musical Objects and flexible metadata shells.

Author Keywords

Music production; studio environments; workflow; metadata; ethnography; ethnomethodology; labeling; DMOs.

ACM Classification Keywords

H.5.5. Information interfaces and presentation (e.g., HCI) – Sound and Music Computing; *Methodologies and Techniques*.

1. INTRODUCTION

Global revenue from the music industry in 2014 was US\$14.89 billion, with 46% of that constituting revenue from digital channels [30]. Global digital revenue has grown year on year since 2009 and music is a major force in the burgeoning digital economy. Within the music industry there has been a revolution in digital production technologies. There are already numerous DAWs (Digital Audio Workstations) available to prospective music producers on the market, with ever more appearing on the scene and offering a bewildering array of different features. At the same time these new technologies often have to be integrated with traditional analogue technologies and within environments and established workflows that have been built up over decades. This is especially the case for mainstream professional music production, which has been visibly slow to embrace the digital revolution. Alongside this, it needs to be recognized that music studio production environments are often very complex. They may encompass numerous synchronous and asynchronous workflows and implicate work on the part of numerous individuals. All of these features together make professional music studios

enormously rich domains to study with regard to the interests of CSCW. An added feature of particular moment for design is the fact that they are also rich in metadata that is used for a wide variety of different purposes. The management of this can present parties within the process with a range of thorny issues that have to be dealt with as they arise if the workflow is to progress.

In this paper we use the findings from two ethnographic studies of distinct parts of the music production process in order to examine some of the ways in which issues relating to workflow (the sequential organisation of the process) and metadata (data about the music and the work undertaken upon it as opposed to the music itself) can result in overheads for a variety of parties: producers; engineers; and even the artists themselves. In some cases, problems may arise because there is metadata and resources relevant to the current activity in place but they prove to not be ‘fit for purpose’ and therefore in need of change. In other cases it can turn out that metadata and resources have been lost along the way or were never present in the first place. This can result in creating things again from scratch. Often metadata that may have value for other parts within the process is only seen to be useful *right now* and is then disposed of. We particularly focus on how legacy mixes within the music production workflow may result in overheads for the creation of future mixes down the line. This can necessitate a range of workarounds to get the job done. Nowhere is this effort more visible than in the initial processes undertaken to get set up for creating a mix.

A key message here is that the reasoning applied to the accomplishment of one workflow at one particular part of the process cannot be assumed to hold at other points within the process. In fact, the reasoning applied at different stages may prove to be at odds, causing disruption and the generation of extra work. This is not unique to music production but can arise in other domains where complex workflows involving numerous different individuals have to be realized as well. This makes the work presented in this paper of general value to researchers interested in workflow and metadata across CSCW.

In the closing part of the paper we pose the question: are there better ways of packaging and exchanging musical information between people and across different stages of the workflow? This leads us to argue that there is a need for a richer and more structured way of being able to associate

different metadata with audio content according to task, individual and perhaps even genre. It should also be more straightforward to switch between these or copy and edit between them. In relation to this we outline design options we are now exploring based upon the notions of Digital and Dynamic Musical Objects (DMOs) (cf. [10]) and the development of flexible metadata shells.

2. THE LITERATURE

At a general level there is a wide range of existing literature in the CSCW and HCI communities that is related to the interests we are surfacing here. This paper can be seen to contribute to three different concerns within that literature. The area most directly addressed relates to matters of workflow and metadata and its use. However, to encompass the full extent of this literature is not feasible within the scope of a single conference paper. We therefore restrict ourselves here to some more specifically relevant work that covers similar issues regarding distributed, cross-organizational workflows and the associated management of metadata. We can also be seen to be more specifically addressing the literature that examines how people handle metadata troubles and the approaches that might be taken to ameliorate these. Finally, the paper also offers a detailed contribution to the very scant body of ethnographic studies of music production, especially with regard to the ways in which music producers handle metadata.

2.1 Previous studies of music production

There are not a large number of studies of music production environments in the HCI and CSCW literature. Even looking further afield there are not so many cases where these kinds of settings have been studied ethnographically. Indeed, some authors have attempted to examine why recording studios may present challenges for the accomplishment of ethnographic work [52]. Faced with potential access issues some investigators have resorted instead to interviewing practitioners who use studio environments (see, for instance, [25] & [26]). Another potential strategy that some have adopted is the use of surveys. Pestana [45, 46] undertook an extensive survey of mix engineers to try and identify mixing best practices. However, these do not include specific methods for managing the workflow across extensively distributed settings. Nor do they include ways to tackle arising issues regarding metadata. In a study that is much closer to the perspective we adopt here, though without any specific interest in technology design, workflow or metadata, Horning [27] examines the use of tacit knowledge by recording engineers in recording studios. She also looks at their collaborative relationship with artists as a feature of the artistic process, but does not explicitly address how workflow and metadata issues may arise. Another more concretely ethnographic study was undertaken by Porcello [47], who recorded conversations during the course of actual recording sessions in music studios. However, the aim here was primarily to try to analyze the professional character of the language being used by recording

engineers. A few more focused studies have stepped away from mainstream recording studios and used ethnography to examine the practices of people engaged in creating remixes [5, 34]. Whilst interesting, however, this body of work is not engaged in the same consideration of extensive workflows and the transition of metadata that we address here. Another domain where cooperative work practices around metadata production and use *can* be similarly found is video and TV production. Studies of this domain, however, do not adopt the same focus upon metadata that we do here, even though there are strong similarities with regard to the scale of production environments and their distribution (see [13, 14, 15, 44] for examples of studies in this area). Thus lessons learnt from our own studies may fruitfully be related to these other kinds of media production environments. All in all the studies we are engaged in seem to break new ground in their attempt to encompass the detailed practices whereby music-related workflows are accomplished by professional practitioners working in various different kinds of music studios.

2.2 Workflow and metadata and bridging the gap

Few studies have looked directly at the relationship between workflow and metadata in music production environments and how metadata may be preserved or transformed across the workflow (Wilmering et al, [55] being a notable exception). However, some studies have looked more generally at workflow and metadata in this way. Matthews and Aston [37], for instance, have looked at the issues relating to managing metadata across the workflows involved in the curation of digital heritage archives. As a result they make several design propositions regarding the support of more effective interfaces. In the same domain others have explored the scope for the automated generation of semantic metadata. This is to reduce what is seen to be a heavy manual overhead [49]. Particular interest has been paid to how metadata is preserved across workflows in the work of scientists where data may have to be re-used and curated by a number of different parties [2]. Here there is particular emphasis upon the possible role of automation (see [3, 4]) rather than providing support for contingent interests. This is similarly the case in other domains (see [36] & [42]). Others are more focused upon how to make metadata visible across different workflows in different software packages [11, 23] or repositories [7, 20]. As is discussed in Section 6, maintaining metadata visibility across different software environments is something that is also relevant to the concerns of this paper, where production can be both geographically and organizationally distributed. The broader problem of maintaining metadata consistency and coherence across complex workflows has also been examined by others such as Pellegrini [42, 43].

2.3 The in situ management of metadata troubles

As far as we are aware no other authors have tackled the specific challenges of localized metadata in complex music production workflows. Nor how such challenges are

currently managed. Some studies, however, do look at metadata issues of a similar order. We have already pointed out several where the production of metadata is considered to be an overhead (e.g. [3]). Some authors (e.g. [24, 32, 33]) point to how metadata can be enormously heterogeneous across different repositories. This has resulted in an interest in tackling issues of ‘interoperability’. There is thus a concern with matters such as how to facilitate metadata exchange and how to ‘interpret’ what different bodies of metadata may mean (see also [50]). Pellegrini [42, 43] considers one of the key challenges in the generation of effective metadata to be the tension between its technical articulation and the language used to refer to it in non-expert communities. Meanwhile Huang & Qin [29] extend the issues beyond interoperability to matters such as ‘portability, reusability, manipulability, sufficiency, and modularity’ (see also [21]). Meyernik [38] conducted ethnographic studies of scientists to uncover what he believes to be 4 key issues in handling metadata across various workflows: the problem of who is actually responsible for its creation; the tension between standards and local ad hoc practices; the distribution of metadata knowledge around organisations; and the role metadata might play at different points within different workflows. The second and fourth of these issues resonate particularly strongly with the results we are presenting here, though Meyernik does not expand upon them to any significant degree. Some authors examine another issue we touch upon here: the extent to which metadata generation may lose important contextual information that is critical to making the material it relates to meaningful [18, 54]. A few studies *do* exist that touch upon the use of metadata in relation to musical activities, but none of these can be seen to be of relevance to the issues we are focused on here (see [1], [6] & [35]).

3. THE STUDIES

The work we are drawing upon here was undertaken as part of a far-reaching UK-based research project that looks at bringing semantic web, signal processing and content-derived metadata technologies to bear across the music industry, from end-to-end, to enhance both production and consumption processes. In order to ground the work of the project within real-world practice we have been undertaking a series of detailed ethnographic studies of how artists and music producers actually work across a range of different settings and genres. For the sake of perspicuity we will be focusing here upon two specific studies: A) the in-studio capture of a semi-professional live demo; and B) a professional post-production session involving the creation of a promotional pre-mix to support the financing and release of an album. This enables us to examine how metadata is first created and how its presence may lead to subsequent troubles.

Studio-based professional production can involve a complex division of labour. It is not just bound up with the number of parties working in the moment to create a viable

recorded product. It may also be temporally stretched across different parties working in different environments and using different tools. In many instances initial preparation of materials may happen in people’s homes, rehearsal studios, even in bars and cafes or via online resources such as dropbox. Actual recording of instruments most often happens in established studio environments with dedicated live rooms, though an increasing amount of material is also recorded in people’s homes. Professional post-production may happen in any of a range of studios, including (but increasingly rarely) larger studios similar to ones where the actual recording of instruments takes place. However, it mostly happens in smaller specially crafted studio spaces focused on post-production, or even just on people’s laptops at home. Mastering may happen in different studios again, wherever the necessary equipment has been assembled. Producers and engineers working in professional music production are increasingly mobile and may work in numerous different environments over the course of any one project. Professional recording projects may also be initiated by the artists themselves or by the record labels. This impacts upon budget and can lead to different kinds of resources being used. All of this means that in a professional recording project many different people in different places may end up working with the same core elements, across many different versions, and towards different ends. Each of these parties may have different working preferences and may bring to bear a different set of working practices. To further emphasize the complexity involved, the actual workflow for a professional project may encompass any or all of the following steps: initiation of the project; circulation of relevant materials (this could be written materials or bits and pieces to listen to); rehearsal; locating and booking a studio and supporting musicians; circulating parts for the supporting musicians; doing a first recording session in the studio; creating rough mixes for a variety of reasons, such as giving the band something to take away, informing assessment of the adequacy of what has been captured, and providing a basis for the next mix in the cycle; adapting the rough mix to make a pre-mix; circulating the pre-mix for assessment; adding and refining parts (overdubbing and re-recording); taking previous mixes through to a ‘final’ mix (the mix can remain open to subsequent revision); mastering; promoting; and distributing a release.

Study A involved observation of a music studio production session, which was organised by a semi-professional ‘Retro Rock’ covers band intending to publish song recordings on their website. They hoped to use the recordings to promote and obtain bookings for their live performances. The band negotiated some free studio time and the help of the resident recording engineer at a professional-standard studio facility installed in a university. This was occasionally available to non-students during holiday periods. The work of 6 people was observed overall: a recording engineer who co-produced the venture with a

singer and a rhythm guitarist; a lead guitarist; a bass guitarist (who was also the backing singer); and a drummer. The study captured three inter-related production workflow steps across three days in the studio, each step covering one whole day: i) *Set-Up*, where the recording engineer, singer, and rhythm guitarist set up and tested the musical and recording equipment; ii) *Recording*, where the whole band performed the songs more or less 'live' for the recording engineer to record to multi-track; and iii) *Overdubs*, where the rhythm guitarist, singer and backing singer performed additional instrumentation and vocals, that the recording engineer dubbed onto the initial multi-track recordings.

Study B centred upon the work of a professional music producer and sound engineer involved in the production of an 'authentic' album of covers of original blues songs from the 1950s and 60s for a professional UK-based blues band. 2 days were spent observing the producer working on his own in another professional producer and engineer's state-of-the-art studio to create a pre-mix of the album. The songs had already been recorded in basic form at a previous session in a different commercial studio and by a different engineer. The goal of the pre-mix was to take the rough mix coming out of the first recording session and to work it up into something more representative of the project. This could then be used as a basis for further judgments about what might still need to be done or changed in some way. It was also intended to provide early promotional material for circulation around promoters and potential investors.

Pre-mix is possessed of its own workflow which takes roughly the following form: acquisition of the rough mixes; importing in the rough mixes so that they can be manipulated in the DAW; setting up some basic elements in the DAW (in this case Pro-Tools) that are maintained across the project (e.g. instrument groupings and their attributes); working through the materials track by track to arrive at a 'satisfactory' mix; exporting the materials from the DAW so that they can be played through by the artists and other interested parties; and actually distributing the pre-mix amongst the artists, or to other parties.

Both studies were ethnomethodologically-informed 'design ethnographies' [8]. They focused upon the in situ production of the people involved's working practices and how they are constituted methodologically. In particular, they exposed people's specific rationales and what it actually took to 'get the job done'. Thus all of the presented materials were acquired during the actual accomplishment of work-in-progress rather than through subsequent interview. As presented in section 6, the findings are now being used to inform the preliminary development of specific design prototypes to demonstrate potential solutions to the issues we have uncovered.

4. THE FINDINGS

In the following sections we shall be examining a number of issues to do with how resources are transitioned across the workflow of music production processes, drawing in

detail upon the studies outlined above to illustrate a number of themes that emerged.

4.1 Creating resources in the recording studio

In this first set of observations we shall be using the materials from Study A that were captured over the course of the actual recording of instrumental and vocal parts in a recording studio. The purpose of this section is to look at how various resources first get created to underpin the workflow and the production of metadata. Study B, where similar initial recordings were being used as the basis of post-production, will then be used to illuminate how at other points in the workflow the initial logic of creation may get set aside.

One of the important roles played by the recording engineer during Study A was the configuration of the equipment during the set up activities. Over the course of this he was seen to methodically add meaningful labelling to the hardware and software surrounding his working position in front of the recording console in the control room. This labeling is a particularly important part of metadata assignment in recording sessions. It provides for the ready recognition of things when having to work with multiple features simultaneously. This is a recurrent aspect of recording studio work where lots of audio channels have to be managed together. In Study A the centrepiece of the Control Room was a large 48 channel Neve recording console. Through this the analogue audio signals were inputted, metered, processed, mixed and subsequently listened to through the studio monitor speakers. The recording engineer labelled each of the recording console channels by writing with a marker pen on a strip of masking tape stuck across the tops of the faders on the front edge of the console (see Fig. 1). This enabled him to identify the channel assigned to a sound source and access the controls quickly to adjust the mix of sounds heard in the monitor speakers. It could also be used for troubleshooting or critical listening. The kinds of information noted on these strips fell into two distinct areas:

1. Channels on the right half of the console (Nos. 25-48), were labelled with type and purpose of the audio source at its input e.g. 'GTR 2, SM57' for a (Shure) SM57 microphone capturing Guitar 2 amplifier.
2. Channels on the left half of the console (Nos. 1-24) were labelled with instrument and headphone stereo sub-mixes outputted from Pro Tools. e.g. 'DRUMS'




Figure 1. Channel labeling strip on the recording console

Once a session had finished these strips were removed from the console and stuck to the control room door. Here they were preserved until the conclusion of the recording project in case there was any need to refer back to them. After this they were thrown away.

According to the recording engineer the labeling he engaged in could serve a number of different purposes such as: denoting redundant channels and preventing him from inputting multiple sources to the same signal path during set up; being able to locate and interact with recording equipment controls during recording; serving as a substitute for making notes; and in case he needed to access the physical configuration data again at a later date.

Alongside all of this the recording engineer also typed *digital* text labels into the Pro Tools DAW software displayed adjacent to the recording console. These were duplication of the labelling on the recording console to indicate the inputted audio sources and outputted sub-mixes, aiding the location of software controls. On completing the set-up work, the recording engineer checked and updated the Pro Tools session's labeling and arrangement before saving it as a template. Loading a template allowed the recording engineer to reinstate the original set of labels and settings, but discard the audio and other changes associated with recording a specific song.

An important part of the recording session is the use of headphone mixes. These allow the performers to hear themselves, the rest of the track and the recording engineer while they are recording. Each musician wore a set of headphones fed with a mix created and managed in Pro Tools by the recording engineer in the control room. The mix was constructed according to preferences communicated to him by each musician. In Pro Tools he labelled each separate mix with the relevant musician's name (see Fig. 2).




Figure 2. Pro Tools mix window showing labelling of tracks, faders, etc. (derived from original data but anonymised)

So far we have detailed the purposing of metadata in a textual form. However the audio from the talkback system was also recorded as a means of prospectively supporting the production workflow, much like its textual counterparts. The recording engineer's rationale for this was twofold: firstly, to record and potentially include a roomy ambience in the final mix; secondly, and most saliently, as a record of musical directions, performance critique, and editorial decisions made in conversation during the session. Details here could include which take of a song was considered best and what edits, if any, were needed. The talkback audio then amounts to metadata in a different mode, as it captures the kind of information that is useful to refer to in the post-production editing stage of the workflow. At one point the recording engineer also used the recording of the talkback audio to recover the count-in to a track when recording an overdub.

4.2 Handling legacy and creating metadata in subsequent mixes

In this section we use materials from Study B to illustrate the ways in which the kinds of resource and metadata creation processes visible in Study A can lead to subsequent issues. This is particularly the case when materials that have been recorded at one studio are worked with by a different person in a different studio to accomplish different ends.

At the beginning of his activity the producer/engineer for the blues album project had received all of the materials recorded at a previous studio on a portable hard-drive. They were then uploaded into the local version of Pro Tools to be worked up as a pre-mix. In this excerpt he is inspecting what is there prior to getting underway with the first track:

Producer: Let's just (looking to the right) I don't need to hear that (clicking on 'S' button above fader on channel third from the right

That was his rough mix ... (Also clicking on 'I' button above that) Now something you can do in Pro Tools is you can make a track

what's called inactive (hovering with mouse over bottom of channel – clicking on symbol that displays option 'Make Inactive')

Straight away what is being dealt with here is a legacy of prior work that he considers to be superfluous to the work he is currently engaged in. He therefore wants to set it aside. There is a track present that plays out the overall rough mix. The previous engineer created this to circulate amongst the musicians so that they could get a rough sense of what had been recorded so far in the studio. This is comparable to the rough mix mentioned in Study A. However, the work the producer is now engaged in is expressly designed to replace this rough mix. It therefore has no place within the current session. Rather than deleting it he decides to make the track 'inactive'. This way it is 'no longer using any CPU resources'.

Other elements that are carried through from a prior session are not just rendered inactive but deleted completely:

Producer: I want to get rid of- (moving cursor to select top item) Yeah, let's get rid of all of these (moving cursor across to adjacent channel) ... These are all the headphone feeds from the studio for the musicians. ... Now probably what will happen is we'll go back in the studio (continuing to perform same action on adjacent channels as talking) and C [the original studio engineer] will set them up again, so I'll end up having to- (gesturing to side with hand) but it's just that they- They kind of litter my screen a little.

This deletion of the headphone feeds is especially interesting in view of all the work visible in Study A getting them set up to support interaction between the recording engineer and the musicians in the studio. Note also how the producer acknowledges that this requirement is not going to completely disappear. At a later stage someone is going to have to create them again.

During our observations an immediate problem was uncovered with the set-up for the various tracks. The producer started to work towards building sub-groups when he discovered the tracks he'd selected were not audible. Having checked a number of things, including the obvious volume controls, he discovered the true source of the issue:

Right, okay, so that's a pain in the arse. I need to- He- He's basically what he's done is he's set all these channels up to his workflow in the studio ... They're all going out to something weird called bus three and four. So what I'm going to do is I'm just going to select all the tracks (does so) ... I'm going to send them all out of output one and two

Creating sub-groups is a central part of preparing for a pre-mix. It brings together a number of associated instruments so that they can be managed as an ensemble. One of the key rationales underpinning the creation of sub-groups is the potential advantage of having a smaller number of channels to work with:

It's just a sub-group that's a- that's all it is and then what I could do in my mix later is I could just have a bunch of sub-groups and I don't even need to worry about what's going on in the individual channels...

The assignation of sub-group attributes, in particular, is all about ensuring that particular behaviours are replicated

across all instruments within the group. One of the core attributes is 'muting'. It is essential that you be able to bring the groups in and out of the mix according to need. By ensuring that all members of the group share the 'mute' capacity this can then be simply accomplished by clicking a button that will mute the group or bring it back in. A similar reasoning may also be applied to the sub-group volume.

Of course, arriving at such ensembles involves deciding what instruments each group should actually contain. Here again metadata issues may arise. In the following extract the producer has just set up the sub-group for the drums when he notices the presence of claves in the mix:

So I've done that for the drums. I'm probably going to go through and do a similar thing- The clave! The cla-vay. What shall we shall we do with that? Let's send that to the drums as well actually, because that's a kind of percussion isn't it

Clearly, in the original recording session the engineer had not thought to group together the claves with the drums in any way. But for the purposes of accomplishing the work he is engaged in now the producer needs to have as few sub-groups as possible. To do this he needs to put in one block anything he might count as percussion.

However, as soon he tries to bring the claves in and mute the group, he discovers he can still hear them. What had happened was that, during the actual recording session in the studio, the claves had been mic'd up in two different ways. One microphone was closer to the instrument and capturing it directly. The other microphone was further away and capturing it together with the ambience of the room. When it came to doing the rough mix the engineer in the studio, for reasons that are not clear to the producer down the line, decided not to use one of these tracks. So he hid it. The work undertaken to eliminate the issue caused by this shows that the logic applied by the engineer when creating a rough mix in the studio, and the logic being applied now to the creation of a pre-mix, are two distinct and not necessarily commensurate things. This resulted in a significant overhead for the producer as he worked on what amounted to, for him, repair. Many of his comments along the way made it clear that, by simply trying to hide the track away rather than making its status evident in the assigned metadata, the recording engineer had given the producer an uphill task to recover the sense of what had been done or how best to repair it.

Another important job with regard to the initial set up for a mix involves the preparation of visual cues that can facilitate rapid navigation amongst the elements. A common practice here is the assignation of colours to different 'types' of element so that their nature can be recognised at a glance. Something that we will see in the interactions with the system that unfolded across the set-up workflow here is that there are several ways of reasoning about what a core 'type' of element might be. At the beginning the reasoning centres upon assigning colours to instrument groupings, such as drums, bass, guitars, and

vocals. However, by the end all sub-group tracks are made one colour and all the other tracks another.

Note here that the various instrument tracks had already been assigned colours in the first session in the recording studio. These colours are preserved as part of the metadata for the files the producer has now uploaded into the local version of Pro Tools. This generates extra work. Right from the outset the producer registers a dislike for the current colour assignments:

(Producer looking at sliders on mixer in Pro Tools) Track colours. I hate the track colours ... This is C [engineer in studio for first session]'s idea. -

In the following excerpt the producer begins to engage in the fine-grained work of actually getting colours assigned to different tracks in a way that suits his own preferences:

Er:m, The other thing I might decide to do at this point is all of the tracks in the group- He's- He's ... used a colouring system that happens to work for him. I probably wouldn't work this way (selecting Preferences from Pro Tools menu) I tend to work this way ... I would set my colours ... probably to groups.

This firmly underscores the fact that the choice of colours as a navigation resource is something that is essentially down to individual preference. At the same time it makes clear that the choice of colours is also doing some kind of work rather than being simply random. At this point the producer started to wonder whether he should leave the track he had created for the sub-group for the drums the same colour as the drums themselves. This led to him colouring the drum sub-group green instead of blue, a first step towards colouring tracks by 'type' instead of by instrument. In discussion the producer acknowledged that he had a notion of standard colours that he liked to use across all of his projects:

See here's an interesting thing, for me I- I often will go for erm blue for bass and I'll do colour associations like that Errm Green for guitars Er Red for vocals because I don't want to kind of- It's The- There isn't any sort of erm (.) There isn't a set way I do this but one thing that I- I do want to get across is that colours are essential for- ... This is an essential part of being able to get around my session And knowing where- ... It's a navigation resource ...

However, within the project in-hand it became apparent that, whilst there was a working sense of the right colours for the right groups in play, the resources offered in Pro-Tools did not simply allow for this to be accomplished as a matter of course within each session he was working upon. Instead he had to work the colour palette to bring about something that approximated to his standard logic, rather than the standard logic being something that he could simply (and 100% accurately) apply. Furthermore, he then began to change strategy and make all of the sub groups a different colour to the mains, whilst retaining a colour that was consistent across all of the subs. This, in turn, led to him deciding that right now it would be even better if he coloured the tracks by 'track type' instead. Where all of this took him was to having all of the sub tracks green and all of the main tracks blue. This may seem to have left him with a

cruder set of resources than the way he was originally heading. However, remember here that his goal is to have a minimal number of faders to move for the larger part of the mixing exercise. His main concern is to get the sub tracks to stand out so that he can use these to drive much of the subsequent activity.

5. DISCUSSION

5.1 The work of managing legacy resources and metadata

Within the above observations we can see how studio engineers and producers need to bring a number of workarounds into play in order to make up for various issues, including metadata conflicts, across the various steps in the process.

In this paper we have focused in particular on the activities involved in setting up for a session. It is certainly the case that troubles can occur at any point in a workflow, and that one of the sources of such trouble can be work others have undertaken before you. However, it was evident to us that much of the potential overhead arising from this is managed by doing set-up work before any other activities are undertaken. This allows for the discovery of elements that may need to be undone or redone to suit local requirements. It also means that any disruption caused by legacy and metadata troubles is positioned so that it does not unduly interrupt the creative process. We are going to look a little more closely here at the work of set-up and four of the primary orientations in play within the process.

5.1.1 *The creation of resources to enable the conduct of the work itself*

Across both of the settings we saw the creation of various labels and other visual affordances. These could either facilitate the rapid navigation of a highly complex environment, or render its complexity more tractable. A key point to grasp here is that the reasoning applied to the creation of such resources cannot be divorced from just where the activity is happening within the workflow, or what it's understood that its outcome should be.

So, in Study B the primary concern was to move beyond what was done initially in the studio and make the product more representative of what the project is understood to be about. Something we saw being undertaken to facilitate the accomplishment of this goal was the aggregation of tracks down in to sub-groups. One might think that pushing everything into sub-groups would make sense for the rough mix as well. However, for a recording studio engineer this is an added overhead. Why would he devote that much more effort to it when his job is nearly done? For the bulk of the work in the studio focusing on the individual instruments makes perfect sense. This is exactly what we saw in Study A.

Another matter that lurks within all of this is that the producer in Study B was also interested in creating a project that sounded 'authentic' for a bunch of covers from the 50s

and 60s. It is part of his knowledge and his competence that he is aware of the limited resources recording engineers and producers had to work with at that time. Recording was bounced to a very small number of tracks. By working with a small number of tracks himself, he is effectively obliging himself to listen to the music and reason about the music in ways more closely connected to the period he is attempting to emulate. This makes it especially clear that, beyond even the specific task being undertaken within the workflow, matters such as genre can have a clear impact upon the construction of appropriate metadata.

Some of the core observations in Study B relate to the use of colouring. Our reason for looking closely at this activity is that it articulates a much broader concern: how to get the metadata to do the right job of work for one's current purposes, rather than to suit purposes elsewhere in the workflow that no longer have any relevance to you. So, within the colouring activity one can actually see a number of important matters regarding how moments within the workflow get reasoned about and accomplished. Key amongst these is the way in which engineers and producers orient to the materials they are working with. Note also that the exercise of different reasoning at different parts within the workflow can result in radically different ways of seeing the tracks, their attributes, and the visual manifestations of their metadata within the DAW in the first place. This can be the case even when the same basic resources and sounds are being worked with. This makes generalisation of features across the workflow, or the creation of technologies that encapsulate such a view, problematic. Furthermore, as we saw in Study B, much time can be invested in a studio in undoing what was done before so that you can re-mould the resources to better fit the work you now need to do.

Something else that was made visible within the colouring activity was the way in which the producer's reasoning about the right strategy to use for it changed over time. What this rather neatly exposes is the way in which just what appropriate metadata should look like is something that can only be uncovered in the process of doing that work and creating that metadata. This is something that strongly challenges any sanguine hope that metadata, especially in the form of labels that are applied and used in the work of production, might take on a generic form that could span the process.

The work of arriving at coherent and locally relevant labelling was another feature of both studies. A key implication of what we saw was that just what components should appear in a sub-group is not something that is given, open to standardisation across different projects, across a whole workflow, or simply provided for by the preceding work in the studio and the creation of the preceding mix. Instead, what counts as a logical member of one particular grouping or another is something that has to be uncovered at just the point you are at now: e.g. that claves should be

lumped together with the drums during pre-mix. This makes it clear that the reasoning being exercised in assigning labels is genre, project and situationally specific. Thus, across the entire workflow naming and grouping metadata may change to suit local logics in play. Thus appropriate naming and re-grouping may be recurrent across larger production workflows.

Another point of potential interest here is the way in which certain kinds of metadata created during various parts of the music production workflow may be treated as something that is tied to a quite specific need and of no further use once that need has been fulfilled (e.g. the handwritten console strips in Study A). A core problem here is that, whilst preserving metadata may on some occasions lead to legacy problems, just disposing of metadata may cause just as many issues. We saw, for instance, the trouble the producer had in unraveling what was happening with the claves tracks as a consequence of one of them being hidden. Nor is it clear that local assumptions regarding whether things may or may not be relevant for future need are always well-founded. The headphone feeds were deleted out of hand in Study B. However, in Study A these were in part recorded to *aid* post-production. Thus metadata that is currently taken to be ephemeral may actually be a key resource for arriving at the intelligibility of production decisions and production practices.

As things stand at present it seems inevitable that, if two conjoint parts of a workflow are working with the same materials but under the auspices of quite distinct understandings of what one is trying to accomplish, troubles will arise, especially if metadata that might help to disambiguate a prior party's reasoning is removed.

5.1.2 The conduct of work that will enable the workflow itself to proceed more efficiently

Something we encountered in Study B was that a good measure of the metadata the producer needed to manage the project was already present in the materials he was given. However, it was in a form that was premised upon the need to a) record the musicians in the studio, and b) create the rough mix he was now expressly setting aside. In that case, for the sake of efficiency, features such as the headphone feeds were simply got rid of and other features, such as the rough mix, were rendered inactive. Thus they could not interfere in any way with the current job of work.

In the recording session the headphone mixes were a key resource. They underpinned much of the social interaction between the musicians and recording engineer. Thus they were necessary to getting the job done effectively. When a recording engineer goes about creating a rough mix of what was captured in a session he has no particular reason to get rid of these feeds. But pre-mix, it turns out, has no business to conduct with headphone feeds. There are no musicians present for the output from the desk to be fed to. So, for the person doing the pre-mix these are considered surplus to requirements and can be removed, together with any

associated metadata. However, it is worth noting that, even as he is getting rid of them, he is commenting that the engineer will doubtless need to create them again the next time they go into the studio to do the overdubbing. This is because at that point they will need to do work as a resource for interaction once again.

One of the key things we have sought to demonstrate in our findings is that handling mixing as part of a larger-scale music production process presents some challenges regarding how to best adapt metadata to suit the purposes of the work you are engaged in now. When existing metadata has to be changed to fit a new logic in play, just what will need to be changed and when it will need to be dealt with is not something that is always predictable. Instead, legacy metadata is uncovered throughout the mixing process. On top of all this, once a problem with labeling etc. has been identified it cannot simply be put off as a minor issue to be dealt with later. Instead making it right becomes something that has to be dealt with as each issue arises. This potentially fragments and disrupts the workflow and generates additional, unforeseen overheads.

5.1.3 The conduct of work that will offset risk and minimize the impact of failure

We have not focused to any great extent in this paper upon the orientations engineers exhibit towards offsetting risk during production. It is, in fact, a significant aspect of the work for them and leads to frequent versioning of tracks and concomitant replication of files and metadata. Primarily the concern is to ensure that actions you are undertaking now will not cause damage to anything that was there previously. However, it should be pointed out that the decision we saw to render tracks inactive rather than dispose of them attests to this same underlying orientation towards not getting rid of things that may subsequently be useful. The interesting thing about this is that there are other features, such as the headphone mixes or the console strips, where there is an assumption that it *is* okay to dispose of them. Here a *local* understanding of the work would suggest that there is no circumstance in which they may be re-used. This obviously has significant implications regarding the role played by individual parties in relation to what kind of metadata *does* get passed down the line.

5.1.4 The verification of available resources and, where necessary, their revision

Wrapped into the work of set-up at numerous points and quite clearly present in many of the examples cited, is a strong orientation amongst engineers towards playing things out. By doing this the consequences of one's actions can be tested aurally, and the current status of materials can be tested. Across the work of both creating new resources and assessing the viability of what is already there it can be seen that, time and again, the actual need to create or modify is uncovered in this fashion. Thus this orientation may be said to underpin all of the other work we have discussed. Unsurprisingly in an aural medium, the ear is the final arbiter. This being the case, the role of the ear cannot

be disregarded, even if the representation of metadata within the DAW is ultimately visual. Metadata is not just *seen* to be not fit to purpose or to need to be modified, it is frequently *heard* to be an issue as well.

5.2 Characterising the necessary workarounds

Of course, parties involved in music production processes cannot necessarily predict just what kinds of metadata fixes they may need to engage in prior to the fact. What *is* provided for in the work we have been describing here is the possibility for key issues to be systematically *uncovered* within the work of set-up. Prior to their discovery within this process there is no certainty: a) that they will encounter redundant elements; b) that they will need to change track colours; c) that they will need to change or create new track attributes; d) what groupings will make most sense on this occasion; e) what metadata might need to be preserved for future use; and so on. The workarounds present in set-up are systematic in the following kinds of ways:

- They provide for the creation of elements that are clearly missing (including metadata of various kinds);
- They provide for getting rid of stuff that clearly doesn't belong (once again, including metadata);
- They provide for the verification of what is there and its revision should this prove necessary (which yet again includes metadata).

Thus the work of set-up in particular merits the kind of close inspection we have been giving it here because it may be instructive for what may offer the most effective forms of future support.

What we have seen in this section is a recurrent need to create metadata across both settings. This need is tightly bound up with the local accomplishment of the work. Much of the metadata established in these circumstances is retained in one way or another as song files move on through the production and post-production process. At the same time much of it is so tightly bound to the specific needs of specific parties at specific points in the workflow that its relevance is lost once it moves on to the next phase. As we saw in Study B, when this is the case it will end up being set aside or revised in order to fit with the new body of requirements. At the same time, as we saw in Study A, some metadata created along the way is so bound up with the local situation that even the parties creating it can see no reason for it to be preserved. This poses a serious question regarding *where* in the process managing metadata and legacy issues should be situated. It also poses a question as to *whose* job it should be to deal with them. At present this happens primarily in set-up, because this is where troubles are most often uncovered. However, if one looks to the heart of the issue it becomes clear that the key moments of significance are the transition points. Here old relevancies are set aside and new relevancies come into play.

6. LOOKING TO DESIGN

The key design implication that we draw from the above is the need for better ways of packaging and exchanging musical information between people and across stages of a workflow. It is, of course, the case that there are already ways of doing this in DAWs. Tracks are routinely exchanged as digital files using encoding standards such as MPEG-3 and MPEG-4, accompanied by the descriptive metadata specified in MPEG-7. Many DAWs also provide facilities for exporting and importing entire projects using interchange formats such as the Open Media Framework [28, 41]. However, even when these work (and in practice it tends to be rather hit and miss), one only gets to export the currently active metadata, track labels and so forth, while the retention of legacy metadata is an issue. Furthermore, this in no way addresses the problem outlined above – that the way metadata is structured and labelled in one session will not necessarily match the preferred logic of another.

We therefore introduce a richer and more structured way of associating metadata with audio content according to the individual and task at hand. Particularly important is maintaining multiple perspectives on the metadata that is present and being able to readily switch between these throughout a music production workflow. This is the focus of ongoing, though currently preliminary design activity.

Dynamic music objects

A potential solution lies in the concept of *Digital Music Objects* that has recently emerged from music technology research. Inspired by previous descriptions of Research Objects from the eScience community [9], De Roure has proposed that the same approach might be applied to musical data and collaborations. He defines a Digital Music Object as a structured unit of musical exchange that “enables ease of reuse and remixing of music right through the chain from composition to consumption” [10]. Others have picked up the baton, extending the concept into *Dynamic Music Objects* where a combination of structured audio (e.g., multiple stems) and rich metadata enable music player software to adapt and remix tracks on the fly, enabling mobile and contextual listening experiences [51]. In short, Digital or Dynamic Music Objects (luckily both abbreviated to DMOs) crystallise audio and metadata into a structured form that can be passed from one person/stage of a musical workflow to another.

Powerful as this concept may be, our study suggests that it requires extension in several respects. Our first contribution is to extend DMOs with multiple ‘shells’ of metadata, each of which supports a particular individual performing a particular task as part of a wider workflow. Thus, an individual engineer undertaking a recording session may create and temporarily invoke a metadata shell that maps the DAW interface to the labels, colours, etc., needed right now for the task of recording. This may be distinct from a shell that is subsequently invoked for creating a rough mix and different again from the shell that another engineer

might invoke for recording or mixing (Fig. 3). The key idea here is that these metadata shells can be readily picked-up and set-aside to configure a DAW to a specific task.




Figure 3. A DMO with multiple metadata shells

Supporting local, individual and tacit metadata

There is already a rich history of exploring the nature of metadata within music technology research. Some of this work focuses on the way in which different kinds of metadata may serve different purposes [53]. It is also central to much of the work on music ontologies [51][17]. However, whilst standardised ontologies may provide a baseline of terminology for talking about the ‘things that matter’ such as ‘tracks’, ‘labels’, ‘colours’ and ‘groups’, we argue that is insufficient to capture many of the tacit practices of labeling and description that we have revealed. Rather, we are proposing that, while each metadata shell may be structured according to a standardized ontology, the fine detail of how things are actually labeled and coloured remains local to individuals and activities and so needs to be both switchable and sharable. This directly reflects our findings that different people may label the same track differently and moreover, that the same person may label the same track differently at different times when undertaking different activities. It also adds a new perspective to existing discussions in the metadata literature regarding ‘personalization’ (see [12] & [48]) and poses a distinct challenge to recent suggestions that the appropriate labeling of metadata might be best handled through crowdsourcing (see [19] & [40]).

As a further comment, we note that the use of metadata is arguably most developed within the domain of libraries and repositories, where it is fundamental to the tasks of discovery, management and preservation. In these domains metadata is typically permanent and immutable, and professionally created and curated to avoid ambiguity. A particular contribution of our work here is that the kinds and uses of metadata present in music production are far more local, ephemeral and contingent, and hence require more flexible and individual support.

Extending Digital Audio Workstation software

Finally, we consider how our extended notion of DMOs might be incorporated into the design of future DAWs. Specifically, we propose that metadata shells should be:

Switchable, so that an individual can choose to apply a new one, while hiding the existing one, so as to reconfigure the interface, for example relabeling and regrouping channels.

Editable, so that individuals can create new ones or customise and repurpose existing ones.

Discoverable, so that individuals can learn about other's mappings. This might encourage them to adopt common labels, groupings and so forth where appropriate. This may serve to promote the emergence of common metadata schemes without enforcing them. Discovery might be enabled by contextual recommendations within the DAW interface, such as mouse-overs revealing prior labels. Discoverable metadata may also support the kinds of highly distributed workflows that are commonplace in music production. For instance, knowing who else may be using your metadata down the line may facilitate shaping it accordingly. Comparing mappings across workflows may also support their optimisation and the emergence of local standards.

Transferable, so that multiple shells can be passed along between activities and DAWs within DMOs. Personal metadata shells may be saved in personal profiles. These may then be used to assist configuring any new DAW for people's own use. At the same time it should be possible to hide or even delete mappings (or parts thereof). Hiding shells might render them inactive but still maintain them within the system so that they can be restored later on. This minimizes the risk of losing important information. Permanently deleting shells from a DMO may be desirable to protect professional know-how, especially when DMOs are transferred across organizational boundaries.

We therefore propose significantly extending both the import and export functions of future DAWs. The process of importing a DMO will involve re-contextualising the material and metadata, including expanding, showing, inspecting, cleaning and filtering what is present; remapping, both from someone else's idiosyncratic terms, colours and labels as well as from more 'official' standards; and annotating and qualifying (e.g. assigning local judgments of quality or usefulness to imported elements). Exporting a DMO mirrors the import process by anticipating what will happen next, including: cleaning and filtering (or even hiding sensitive or confidential details); mapping and translating, perhaps to more standard terms, or to anticipate known idiosyncrasies of future users; and recording and declaring hitherto tacit or implicit information (e.g. the identities of musicians, engineers, specific equipment used, and legal matters).

7. CONCLUSION

In this paper we have examined studio music production environments and focused in particular upon how the use of metadata is embedded within the local accomplishment of a complex series of workflows. Existing metadata research would suggest the likelihood of potential issues here. However, to date there is a paucity of ethnographically detailed work that exposes exactly what these issues might look like and how they might ramify. The studies we have

undertaken have uncovered a number of important tensions across workflow steps between various production environments. These particularly relate to the local usability of metadata, its visibility, and its retention for future use. We have also noted that these tensions can generate significant overheads, causing delay and interruption and potentially limiting the scope of what may be accomplished. On the basis of these observations we have put forward some suggestions regarding how metadata may be better articulated across production workflows, with a special emphasis upon the possible role of DMOs and flexible metadata shells. This is likely to be of benefit to other production domains with equally distributed and fragmented workflows, especially where metadata use is similarly open to local variation. Our current work on realizing the design ideas articulated here is ongoing and will be reported in subsequent publications.

ACKNOWLEDGMENTS

This research was supported through the UK EPSRC project Fusing Semantic and Audio Technologies for Intelligent Music Production and Consumption (FAST) (EP/L019981/1). This paper includes all essential supporting data. The respective field workers hold associated field notes and video recordings in confidence.

REFERENCES

1. Ahmed Ahmed, Steve Benford, and Andy Crabtree. 2012. "Digging in the crates: an ethnographic study of DJs' work." *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM.
2. Anastasia Ailamaki, Verena Kantere, and Debabrata Dash. 2010. "Managing scientific data." *Communications of the ACM* 53, no. 6 (2010): 68-78.
3. Pinar Alper, Carole A. Goble, and Khalid Belhajjame. 2013. "On assisting scientific data curation in collection-based dataflows using labels." In *Proceedings of the 8th Workshop on Workflows in Support of Large-Scale Science*, pp. 7-16. ACM.
4. Khalid Belhajjame, Katy Wolstencroft, Oscar Corcho, Tom Oinn, Franck Tanoh, Alan William, and Carole Goble. 2008. "Metadata management in the taverna workflow system." In *Cluster Computing and the Grid, 2008. CCGRID'08. 8th IEEE International Symposium on*, pp. 651-656. IEEE.
5. Phillip D. Brooker and Wes Sharrock. 2013. "Remixing Music Together: The Use and Abuse of Virtual Studio Software as a Hobby." In Peter Tolmie and Mark Rouncefield (eds) *Ethnomethodology at Play*, Ashgate. 135-155.
6. Alan Chamberlain and Andy Crabtree. 2016. "Searching for music: understanding the discovery, acquisition, processing and organization of music in a domestic setting for design". in *Personal and Ubiquitous Computing*. pp. 1-13. Springer

7. George Chin Jr and Carina S. Lansing. 2004. "Capturing and supporting contexts for scientific data sharing via the biological sciences collaboratory." In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, pp. 409-418. ACM.
8. Andrew Crabtree, Mark Rouncefield, and Peter Tolmie. 2012. *Doing design ethnography*. Springer Science & Business Media.
9. David De Roure. 2013. "Towards computational research objects," in *Proceedings of the 1st International Workshop on Digital Preservation of Research Methods and Artefacts co-located at Joint Conference on Digital Libraries*, Indianapolis, 2013.
10. David De Roure. 2014. "Executable music documents," in *Proceedings of the 1st International Workshop on Digital Libraries for Musicology*, 2014, pp. 91-3
11. Ewa Deelman and Ann Chervenak. 2008. "Data management challenges of data-intensive scientific workflows." In *Cluster Computing and the Grid, 2008. CCGRID'08. 8th IEEE International Symposium on*, pp. 687-692. IEEE.
12. Marcos A. Domingues, Camila V. Sundermann, Flávio MM Barros, Marcelo G. Manzato, Maria GC Pimentel, and Solange O. Rezende. 2015. "Applying multi-view based metadata in personalized ranking for recommender systems." In *Proceedings of the 30th Annual ACM Symposium on Applied Computing*, pp. 1105-1107. ACM.
13. Arvid Engstrom, Mattias Esbjornsson, Oskar Juhlin, and Mark Perry. 2008. Producing collaborative video: developing an interactive user experience for mobile tv. In *Proceedings of the 1st international conference on Designing interactive user experiences for TV and video (UXTV '08)*. ACM, New York, NY, USA, 115-124.
14. Arvid Engström, Oskar Juhlin, Mark Perry, and Mathias Broth. 2010. Temporal hybridity: footage with instant replay in real time. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '10)*. ACM, New York, NY, USA, 1495-1504
15. Arvid Engström, Mark Perry, and Oskar Juhlin. 2012. Amateur vision and recreational orientation:: creating live video together. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work (CSCW '12)*. ACM, New York, NY, USA, 651-660.
16. Steve Everett. 2001. "Logics of Value for Digital Music." *Proceedings of the International Conference Music Without Walls? Music Without Instruments*. 2001.
17. György Fazekas and Mark B. Sandler. 2011. The Studio Ontology Framework, 12th International Society for Music Information Retrieval Conference (ISMIR 2011), International Society for Music Information Retrieval
18. Muriel Foulonneau, Timothy W. Cole, Thomas G. Habing, and Sarah L. Shreeves. 2005. "Using collection descriptions to enhance an aggregation of harvested item-level metadata." In *Digital Libraries, 2005. JCDL'05. Proceedings of the 5th ACM/IEEE-CS Joint Conference on*, pp. 32-41. IEEE.
19. Jonathan Fürst, Kaifei Chen, Randy H. Katz, and Philippe Bonnet. 2015. "Demo Abstract: Human-in-the-loop BMS Point Matching and Metadata Labeling with Babel." In *Proceedings of the 2nd ACM International Conference on Embedded Systems for Energy-Efficient Built Environments*, pp. 101-102. ACM.
20. Yolanda Gil, Jihie Kim, Gonzalo Florez, Varun Ratnakar, and Pedro A. González-Calero. 2009. "Workflow matching using semantic metadata." In *Proceedings of the fifth international conference on Knowledge capture*, pp. 121-128. ACM.
21. Sara Graves. 2011. "Meeting the challenges of data-intensive science." In *Proceedings of the 2011 workshop on Climate knowledge discovery*, pp. 4-4. ACM.
22. Kristina Groth and Jeremiah Scholl. 2013. "Coordination in highly-specialized care networks." In *Proceedings of the 2013 conference on Computer supported cooperative work companion*, pp. 143-148. ACM.
23. Andrew Harrison, Ian Harvey, Andrew Jones, David Rogers, and Ian Taylor. 2011. "Object reuse and exchange for publishing and sharing workflows." In *Proceedings of the 6th workshop on Workflows in support of large-scale science*, pp. 67-76. ACM.
24. Bernhard Haslhofer and Wolfgang Klas. 2010. "A survey of techniques for achieving metadata interoperability." *ACM Computing Surveys (CSUR)* 42, no. 2 (2010): 7.
25. Antoine Hennion. 1981. *Les professionnels du disque* (Paris)
26. Antoine Hennion. 1989. "An intermediary between production and consumption: the producer of popular music." *Science, Technology & Human Values* 14, no. 4 (1989): 400-424.
27. Susan Schmidt Horning. 2004. "Engineering the Performance Recording Engineers, Tacit Knowledge and the Art of Controlling Sound." *Social Studies of Science* 34, no. 5 (2004): 703-731.
28. Matt Houghton. 2010. Transferring Projects between Different DAWs, Sound on Sound: Tips and Techniques, December 2010,

www.soundonsound.com/sos/dec10/articles/daw-transfers.htm (verified July 2016)

29. Hong Huang and Jian Qin. 2013. "Understanding metadata functional requirements in genome curation work." *Proceedings of the American Society for Information Science and Technology* 50, no. 1 (2013): 1-4.
30. International Federation of the Phonographic Industry. 2015. Digital Music Report 2015: Charting the Path to Sustainable Growth. IFPI
31. Wenxin Jiang, Xin Zhang, Amanda Cohen, and Zbigniew W. Raś. 2010. "Multiple classifiers for different features in timbre estimation." In *Advances in Intelligent Information Systems*, pp. 335-356. Springer Berlin Heidelberg.
32. Michael Khoo and Catherine Hall. 2010. "Merging metadata: a sociotechnical study of crosswalking and interoperability." In *Proceedings of the 10th annual joint conference on Digital libraries*, pp. 361-364. ACM.
33. Sébastien Laborie, Ana-Maria Manzat, and Florence Sèdes. 2010. "A generic framework for the integration of heterogeneous metadata standards into a multimedia information retrieval system." In *Adaptivity, Personalization and Fusion of Heterogeneous Information*, pp. 80-83. Le Centre De Hautes Etudes Internationales D'informatique Documentaire.
34. Brett Lashua. 2006. "The arts of the remix: ethnography and rap." *Anthropology Matters* 8, no. 2 (2006).
35. Jessica Lingel. 2012. "We realized we had to become librarians: DJs, information practices and music libraries." In *Proceedings of the 2012 iConference*, pp. 569-571. ACM.
36. Philip Maechling, Hans Chalupsky, Maureen Dougherty, Ewa Deelman, Yolanda Gil, Sridhar Gullapalli, Vipin Gupta et al. 2005. "Simplifying construction of complex workflows for non-expert users of the Southern California Earthquake Center Community Modeling Environment." *ACM SIGMOD Record* 34, no. 3 (2005): 24-30.
37. Paul Matthews and Judith Aston. 2012. "Interactive multimedia ethnography: Archiving workflow, interface aesthetics and metadata." *Journal on Computing and Cultural Heritage (JOCCH)* 5, no. 4 (2012): 14.
38. Matthew S. Mayernik. 2010. "Metadata tensions: A case study of library principles vs. everyday scientific data practices." *Proceedings of the American Society for Information Science and Technology* 47, no. 1 (2010): 1-2.
39. Paolo Nesi and Marius B. Spinu. 2005. "Music Score Watermarking." *IGI*, 2005: 2074-2079.
40. Johan Oomen and Lora Aroyo. 2011. "Crowdsourcing in the cultural heritage domain: opportunities and challenges." In *Proceedings of the 5th International Conference on Communities and Technologies*, pp. 138-149. ACM.
41. Open Media Framework, OMF Interchange Specification Version 2.1, Avid Technology Inc, September 1997, www.linuxmedialabs.com/Downloads/LSI/omfspec21.pdf (verified July 2016).
42. Tassilo Pellegrini. 2012a. "Semantic metadata in the news production process: achievements and challenges." In *Proceeding of the 16th international academic mindtrek conference*, pp. 125-133. ACM.
43. Tassilo Pellegrini. 2012b. "Integrating linked data into the content value chain: a review of news-related standards, methodologies and licensing requirements." In *Proceedings of the 8th International Conference on Semantic Systems*, pp. 94-102. ACM.
44. Mark Perry, Oskar Juhlin, Mattias Esbjörnsson, and Arvid Engström. 2009. Lean collaboration through video gestures: co-ordinating the production of live televised sport. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '09)*. ACM, New York, NY, USA, 2279-2288
45. Pedro Duarte Leal Gomes Pestana. 2013. "Automatic mixing systems using adaptive digital audio effects." PhD diss., Universidade Católica Portuguesa.
46. Pedro Pestana and Joshua Reiss. 2014. "Intelligent audio production strategies informed by best practices." In *Audio Engineering Society Conference: 53rd International Conference: Semantic Audio*. Audio Engineering Society.
47. Thomas Porcello. 2004. "Speaking of Sound Language and the Professionalization of Sound-Recording Engineers." *Social Studies of Science* 34, no. 5 (2004): 733-758.
48. Melike Sah, and Vincent Wade. "2011. Automatic mining of cognitive metadata using fuzzy inference." In *Proceedings of the 22nd ACM conference on Hypertext and hypermedia*, pp. 37-46. ACM.
49. Hendrik Schöneberg, Hans-Günter Schmidt, and Winfried Höhn. 2013. "A scalable, distributed and dynamic workflow system for digitization processes." In *Proceedings of the 13th ACM/IEEE-CS joint conference on Digital libraries*, pp. 359-362. ACM.
50. Vishal Sikka. 2005. "Data and metadata management in service-oriented architectures: some open challenges." In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, pp. 849-850. ACM.
51. Florian Thalmann, Alfonso Perez Carrillo, György Fazekas, Geraint A. Wiggins, and Mark Sandler. 2016.

- "The Mobile Audio Ontology: Experiencing Dynamic Music Objects on Mobile Devices", IEEE Tenth International Conference on Semantic Computing (ICSC '16), pp. 47-54, IEEE
52. Paul Thompson and Brett Lashua. 2014. "Getting It on Record Issues and Strategies for Ethnographic Practice in Recording Studios." *Journal of Contemporary Ethnography* (2014)
53. Gabriel Vigliensoni, John Ashley Burgoyne, and Ichiro Fujinaga. 2013. "Musicbrainz for The World: The Chilean Experience." In *ISMIR 2013*, pp. 131-136.
54. Karen M. Wickett and Allen H. Renear. 2012. "The logical form of the proposition expressed by a metadata record." In *Proceedings of the 12th ACM/IEEE-CS joint conference on Digital Libraries*, pp. 413-414. ACM.
55. Thomas Wilmering,, Kevin Page, György Fazekas, Simon Dixon, and Sean Bechhofer. 2015. "Automating Annotation of Media with Linked Data Workflows." In *Proceedings of the 24th International Conference on World Wide Web Companion*, pp. 737-738. International World Wide Web Conferences Steering Committee