

ORGANIZING DIGITAL MUSIC FOR USE: AN EXAMINATION OF PERSONAL MUSIC COLLECTIONS

Sally Jo Cunningham , Matt Jones, Steve Jones
Department of Computer Science
University of Waikato
Hamilton, New Zealand

ABSTRACT

Current research on music information retrieval and music digital libraries focuses on providing access to huge, public music collections. In this paper we consider a different, but related, problem: supporting an individual in maintaining and using a personal music collection. We analyze organization and access techniques used to manage personal music collections (primarily CDs and MP3 files), and from these behaviors, to suggest user behaviors that should be supported in a personal music digital library (that is, a digital library of an individual's personal music collection).

1. INTRODUCTION

The music retrieval/digital libraries literature has focused on the problems of supporting large scale, public digital libraries, nearly to the exclusion of considering how individuals might organize and access personal collections. This is a surprising omission, considering the popularity of the digital CD and MP3 format; even a child may now have a sizeable digital music collection, of a size that requires more in the way of access support than a simple list of filenames or song titles.

What features or functions should a music digital library system include, if it is intended to support individuals in accessing and managing their own music? We search for clues to answer this question by analyzing the ways that people currently access and organize their personal music collections. Insight into everyday music-related activities can have practical implications for design of a personal music digital library—that is, a collection of an individual's music documents, owned, 'used', and organized by that person.

This paper is organized as follows: Section 2 discusses previous research in eliciting music information seeking behaviors; Section 3 describes the methodology used in this paper; Section 4 presents the observed music

implications for the design of a personal music digital library; and Section 5 summarizes this work.

2. PREVIOUS WORK

At present, there is a dearth of a-priori research on music information behavior. Much of the existing music information retrieval and music digital library research has been technology-driven, and music digital libraries as reported in the research literature are largely developed as proof-of-concept demonstrations of the potential of a given tool or effectiveness of a retrieval algorithm, or are focused around providing access to an available set of music documents [8]. Current efforts at studying MIR system usability focus on user behavior exhibited in specific MIR systems, for example by examining transaction logs [13]. While usability studies can suggest improvements to existing software, they are impoverished sources of knowledge about additional features that might be useful or other information behaviors that could be supported.

Research examining human perception and cognition of music is primarily focused on problems in creating software that can match music elements or extract musical phrases in such a way as to produce retrieval results acceptable to human users [2], or on the factors that may influence a user in creating effective music queries (for example, in creating a good 'sung' query to a query-by-humming interface [15]).

A previous, large-scale ethnography of music behavior [4] presents interviews with 41 participants; the focus was primarily on the emotional relationships that participants had with their music, and to a lesser extent on the ways that people use music in their daily lives. These interviews are 'raw' ethnographies—that is, the data is minimally edited for presentation but is not analyzed to induce a theory or explanation of the self-reported behaviors. This study is not of direct use in suggesting design considerations for a music digital library, but could be mined for evidence of music-related activities.

There is a small but growing body of work on music behavior of non-specialists (that is, people who are interacting with music primarily for personal pleasure rather than professionally). The goal of this research is primarily to develop an understanding of how to effectively support access to public collections. In [12], the terminology that participants use to characterize classical music pieces is analyzed and contrasted with formal

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page.

© 2004 Universitat Pompeu Fabra.

organization and usage behaviors, and discusses their

(bibliographic) descriptors. Music queries posted to a music-focused Usenet newsgroup [7] and to the GoogleAnswers ‘ask an expert’ service [1] provide clues as to the types of music documents that people may be interested in obtaining from a public music digital library and the attributes that people can provide to describe their music information need. The strategies that people natively employ in searching and browsing for music in CD stores and public libraries are detailed in [6].

3. METHODOLOGY

The data gathering techniques employed in this investigation were ‘personal ethnographies’, interviews and on-site observations of personal music collections, observations of music store layout and shopping behaviors, and focus groups.

The bulk of the data was gathered through a project assigned to students in a third year university human-computer interaction course. The students were directed to perform a ‘personal ethnography’, in which they examined their own music collections and created a description of their collection’s organization, the collection’s contents, when and under what circumstances they use the collection, and the ways in which they access the collection (e.g., listening to songs, loaning music to friends, reading CD inserts, and so forth). In a personal ethnography or autoethnography [5], ethnographic techniques of observation and analysis are applied to one’s own experiences; the challenge is to view oneself objectively, to see one’s own worldview as freshly as possible and to then interpret the identified experiences in the light of applicable theory.

The students then performed a similar ethnographic observation of a friend’s collection and interviewed the friend to clarify the organizational principles that the friend used in his/her music collection, and to create a description of how and when that friend uses the music collection.

In total, the students conducted ten personal ethnographies and ten observations/interviews of another. The researchers performed an additional six observations/interviews focused on personal music collections. We also draw on eight interviews on music behavior conducted to support an earlier study [6].

Each of the ten students also examined at least one music store’s layout, and performed participant observations of shoppers in the stores. In an earlier paper [6] we argue that as CD stores are a common source of music for many people, the searching and browsing strategies observable in these stores can be useful sources of data on music information behaviors. In this current study, the music store observations provided a commercial view of how music can be organized to facilitate access (although in this case ease of access is confounded with impetus to purchase), and additional data on how people navigate a large, public music collection.

Three focus groups (one of six individuals, two with three participants) were also organized by three of the

students, to solicit experiences with current music organization/playing systems such as MP3 players, and to brainstorm ideas on the functions and features that an ideal music system would include.

Approximately 120 pages of data were gathered from all sources. The data was analyzed using a grounded theory approach [9]. With this technique researchers attempt to approach the data without prior assumptions, and to generate theory from the data. Further qualitative studies or quantitative experiments can then test the validity of the emergent theory. The aim here is to describe how people currently organize their music collections, to suggest features and functions that should be included in a personal music digital library software system.

4. HOW PERSONAL MUSIC COLLECTIONS ARE ORGANIZED AND ACCESSED

The following sub-sections summarize the characteristics of music collections and the observed ways that participants organize, search, browse, and use their personal music collections.

4.1. Collections vary in size and media

The collections varied widely in size—from a single CD owned by an eleven-year old girl (“I don’t get any allowance at all! I can never ever ever afford anything like CDs.”) to an estimated seven hundred plus CDs accumulated over more than a decade. Surprisingly, the organizational schemes employed were relatively consistent over a range of sizes; once more than a handful of music had accumulated (Sections 4.4 & 4.5).

Music collections included a variety of media: primarily CDs, frequently MP3s (ripped from CDs, emailed from friends, or downloaded from the Internet), and older formats such as cassettes, eight track tapes, and vinyl LPs (with older formats seldom or never accessed; see Section 4.5).

At present, music is usually legally obtained as an ‘album’—a collection of songs, as released by the artist. Albums may be burned to CD for use as backups, and compilation CDs may be created, composed of favorite songs from different albums. These physical CDs (purchased or burned) are then physically organized, and the organization and access of these physical collections is the focus of much of this paper.

The organization of individual songs is becoming a more significant activity in personal collections, as the ready availability (both legally, through online music stores such as iTunes or Rhapsody, and illegally, through music sharing services) of single tracks entices music lovers to obtain only the specific songs in a collection that they like. The creation of playlists and compilations [Section 4.6] suggests that even if the album truly does die [3], users will still wish to manipulate groups of songs as well as to access individual tunes.

4.2. Collections are distributed

Few participants with music collections of more than a nominal size (say, more than 20 CDs) keep their entire collection in one physical spot. Collections are generally divided into the active items (that is, those that see regular or occasional use) and the archival items (music that is seldom or never listened to; for further discussion, see Section 4.4).

The active set is also frequently divided into several sub-collections: a small set of very frequently used music, generally placed on top of the CD player or by the computer (see Section 4.2); a large set of occasionally listened to CDs (see Section 4.3) in a CD tower, drawer, or cabinet, beneath or near the main listening device (usually a stereo, occasionally a computer); a set of CDs in a CD wallet, that are played in more than one location and so need to be easily transportable (“[my] CD wallet ... is usually situated either underneath my Discman, on top of the stereo or in the car...”); a set of CDs stored at the workplace or at a university computer lab; and if the home contains several music playing devices, a set of CDs may be associated with each (for example, CDs in the family room, CDs in an individual’s bedroom, and CDs beside the home computer). A degree of forward planning is required to ensure that the right CDs are in their correct locations for listening.

It is clear that this geographic distribution is almost entirely due to the fact that the CD is a physical object, and so must be toted from place to place, if for no other reason than to ‘rip’ it and put the copy on an MP3 player or hard drive. The participants generally viewed having subsets of their collection in more than one spot as an annoyance, generally minor but occasionally major, since locating a desired CD might involve looking in many, sometimes widely separated, places.

Storing a collection in a portable music appliance such as an MP3 player would finesse this problem of geographic distribution, since the entire collection could be then be easily transported to wherever the owner wishes to use it.

4.3. Emergent Structure

A common realization by the students performing the autoethnographies, and indeed by other participants as their interviews proceeded, was that even a seemingly disorganized collection frequently had an implicit structure that had arisen through use:

Before beginning this project I did not think that my collection was organized in any specific way. However after examining it ... though not organized in a traditional way such as alphabetically or by genre. There is a system, which I have implemented without really realizing it.

The most frequently observed emergent structure is a small stack of CDs that are currently receiving a large amount of use. This music is typically located close to the

playing device, to make it as easy as possible to quickly select the CDs for playing. The last played CD is usually placed on top, so that the less played CDs drift to the bottom of the stack. The size of this set of most active music is small—sometimes only three or four CDs, sometimes as many as twenty. Often a limiting factor on size is that a tall stack is prone to accidentally topple, or that it looks messy. When the stack becomes too large, then it is pruned and less frequently played CDs are returned to the main collection.

This small stack organization may be associated with the ‘thrashing’ and ‘sickness’ listening cycle. A few participants reported that a new CD will be thrashed, “i.e. played over and over, until it eventually loses [sic] its novelty”—at which point ‘sickness’ sets in, the CD owner “decides that its [sic] time to listen to something else”, and the CD is moved further from the top of the current stack, or even put with the main set of CDs.

4.4. The Main Active Collection

As noted in Section 4.1, the bulk of most collections are stored near the primary listening device. Where the most frequently listened to items are in a small stack (Section 4.2) and the never listened to items are in storage (Section 4.4), the remainder—indeed the majority of most collections—are only occasionally listened to. A variety of organizations are used for these items:

- by date of purchase, for example with the newest CDs placed either at the top of a stack or at the end of a shelf
- by release or recording date
- by artist, with the artists arranged alphabetically
- by genre, where the number of genres can be large or small (“rap and other”)
- by country of origin (e.g., “New Zealand music”)
- from most favorite to least favorite
- in order of recency in which the CDs have been played

A collection is generally organized into relatively few categories (for example, into very broad genres such as Jazz and Pop). A secondary organization may be applied to each of the broad top-level categories (for example, sorting by artist within genre). The classification scheme is rarely more than two levels deep, so that a linear search is generally needed to locate a particular CD within a category or sub-category. This type of loose ordering is provides acceptable access support, as most collections are small enough that CDs can be located relatively quickly.

Unfortunately, the structure of most collections tends to deteriorate over time. Few people have the patience to return a CD to its proper spot after it is played, guests may disturb a collection while browsing it (Section 4.6), a CD tower may be knocked over and hastily shoved back into place, and so on. Many participants had abandoned a former ordering, and were now simply adding in CDs to the top of a tower as they were purchased (“Once I used to sort by artist, but not any more. It’s too much of a pain.”).

Given the initial interest that is shown in ordering a developing collection, it seems likely that software facilities that will support the organization of music would be welcomed—if, and this is a big ‘if’, the organizational metadata can be quickly and easily added with a new piece of music. If much effort at all is required beyond a couple of clicks, then it seems likely that the metadata tagging will be deferred indefinitely and the digital collection will also subside into disorganization.

4.5. The Archival Collection

Only two of the participants reported that music which had fallen from favor was discarded or allowed to drift off (“they usually end up getting misplaced and lost ... I don’t really pay attention to where I put them.”) Typically if a personal music collection has been accumulated over a significant period of time, then it is likely to include items that are rarely, if ever, listened to. These form the ‘archives’ of a collection, stored away in a closet or otherwise put out of the way. Sometimes these items are archived because the media is out of date (eight track tapes, vinyl LPs), and a player is not available (“approximately seventy vinyl L.P.’s now in permanent storage due to the lack of a turntable (a.k.a. Gramophone)”). Other items are simply no longer of interest to the owner: “music that I have grown out of”.

Why are these items stored, and not discarded? Sometimes it’s simple inertia on the part of the collector—an unwillingness to take the time to sort out the potentially listenable from the completely outgrown. Sometimes the collecting instinct is too strong to resist:

Interviewer: Why are you keeping all of those LPs in your closet?

Husband: I haven’t the faintest idea.

Wife: Because he’s a hoarder!

And sometimes the music is kept because of emotional ties or as a memento: “nothing more than a reminder of changes in my personal taste as I have grown.”

It appears likely that a digital music collection will also eventually include music that the owner no longer wishes to listen to, but is reluctant to delete. An archival facility is likely to be useful—perhaps semi-automatic, with the system suggesting candidate songs that have not been listened to in months or years. A secondary use for this suggestion function would be to remind users about music that they had forgotten about, but that they still might enjoy playing.

4.6. Idiosyncratic Genres: Characterizing Music by Intended Use

One notable way that participants characterized music was by intended use—that is, based on the event or occasion at which they intended to listen to a particular set of music. Music of this type might be listened to as a set of

complete CDs, or might consist of individual songs pulled together into a playlist or compilation CD.

Using the term ‘genre’ loosely, the participants identified a diverse set: programming music (“[techno music] is great to program to, it keeps you typing, even if what you type is nonsense”); detention music (a high school teacher described selecting the music she plays when sitting with students who are serving an in-school suspension: “When I’m working on detention I pull out the Roger Miller tape, the one with the rankest, most country accent and words, and play it for my hip hop kids, so they never want to serve detention with me again”); music to amuse children (“silly songs for the kids, like ‘Please Mr. Custer’, ‘Ahab the Arab’, ‘Transfusion’, ‘They’re coming to take me away, haha’”); driving music (“Everything by Jethro Tull, and one or two others [CDs] that rotate”); “[music that will] keep me awake on a late night drive home from a tiring day on the mountain”); work music (“more ambient music, not as loud and aggressive as some of the other CDs [in the collection]”), mood altering or matching music (“Browsing through the collection to select one that suits my mood, either relaxing if at the end of a difficult day, or something exciting if I am feeling bored”; “[to] cheer me up”; “I only listen to him [artist] when I’ve split up with someone”), and so forth. One of the more fascinating aspects of this study is the sheer number of idiosyncratic genres that emerge from the interviews and observations.

Note that the criteria defining the music intended for a particular use vary—in the above definitions, programming music includes a well known genre (techno) and detention music is selected as being the antithesis of hip hop; the first definition of driving music and detention music are identified more or less closely with a particular artist; silly songs for kids have amusing, G-rated lyrics and a sing-able, simple melody; work music is soft, not “aggressive”, and is used as background noise rather than closely attended to, while the second type of driving music is loud and fast-paced, to keep a sleepy motorist awake; and mood music is may be dependent on any number of facets, including personal associations with events experienced while a particular song happened to be playing (think, for example, of a couple identifying “our song” with a romantic mood).

A facility to allow a user to create personal genres and to easily add metadata to identify music in these genres would be useful in a music digital library. This would be particularly useful as new music is added to a collection with an existing set of user-defined categories. It is easy to imagine, however, circumstances in which the collection owner will miss the opportunity to tag a song with its appropriate genre—for example, if a new genre is being added to a large existing collection. When defining a new personal genre, the individual has at hand exemplars of that genre; locating additional candidates for that genre could be supported by a facility that searches within the collection for ‘more songs like these’. A next step is to clearly identify the musical facets most useful for

characterizing genres—timbre, instrumentation, rhythm, etc—and to develop interfaces for specifying musical query-by-example searches. Research into techniques to automate the creation of personalized playlists (for example, by automatically retrieving and ordering songs with features similar to a user-selected ‘seed’ song) shows promise in this direction (for example, [16])

Music grouped into such a personal genre may be copied onto one or more compilation CDs. This music may be played sequentially by track, if the user has a strong sense how a mood may be developed through a particular ordering of songs, or the player may be set to play the songs in random order. Random ordering can add a sense of variety and novelty to a playlist.

4.7. Collections May Be Shared

While individuals have their personal music collections, they may also participate in a shared collection with others—for example, students sharing accommodations may keep a stack of CDs by the living room stereo, or families may have developed a shared collection that everyone can contribute to and play.

A major drawback experienced with shared CD collections is that they are even more difficult to keep in an intelligible order than individual music collections. The emergent ‘current favorites’ stack organization fails when more than one person is involved; for example, in a family of seven, the current listening stack by the computer consists of “the favourite albums of various family members and [the stack] is in random order as each user searches through the stack until they find their current favourite disk and return it to the top of the stack when finished.” Again, this is a problem tied to the physicality of CDs; different people could view a set of MP3 files in different orderings.

Another common form of sharing occurs when guests are invited to browse a collection to select music to be played during their visit. The music is then part of the social occasion, listened to together or collectively unattended as background to a party. Browsing a friend’s music collection may provide an opportunity to learn more about a new type of music, or to re-think aspects of one’s own tastes. One student, for example, reported that after examining a friend’s collection he re-organized part of his own collection according to the distinction his friend made between New Zealand and international artists.

Not all music lovers are comfortable allowing others to browse or access their collections. One notable exception was James, who had the most elaborate and well-maintained organization for his extensive set of CDs:

...James is adverse to other people selecting CD’s from his collection. For this reason he keeps his collection in his bedroom to “restrict” access to others. On the odd occasion, for example during a party, that his collection is interfered with and the logic disrupted, James will spend time restoring the stand to a state that

is as close as possible to how it was organized before the disruption occurred.

Perhaps if it were easier to share music and to browse the collection without running the risk of disturbing its structure, collectors such as James would be less averse to exposing their music to others.

Another reason cited for reluctance to allow others to browse a personal collection is self-consciousness about one’s musical tastes: “My collection also contains ... CDs I sometimes play but am embarrassed to possess (see: Chris Isaak).” One participant even organized his CDs in a set of racks so as to allow him to hide some of his music: “I can rotate my rack in a way that “shows off” my best CDs while partially obscuring the average and embarrassing CDs.” This participant most eloquently expressed the relationship that a music collection can have with a person’s image:

... I feel that my character is partially judged on the contents of my collection, as I myself consider the contents of a person’s music collection when evaluating what type of person they are. From that last point, I can conclude that an important factor of my music collection [is] that it displays prominently the better/brighter aspects of my personality (see David Grey, Coldplay); while partially obscuring the darker side (see Nine Inch Nails, Deftones).

The ability to customize what others see of one’s collection, and how it appears, may be a crucial feature to some.

While allowing others to browse and listen to one’s music is generally enjoyable, actually loaning a physical CD is generally avoided: “...I do not lend out CDs, as I manage to lose and damage them quite well on my own.” This problem would not exist, of course, if a collection was entirely held on a computer or music appliance that supported easy copying to other digital media—but it is difficult to imagine that such sharing would be legal in the near future. In the meantime, those who have copied their CDs are generally willing to loan the copies (although not the originals). Loaning may be seen as a more significant act than simply handing over a bit of plastic, as it involves a sharing of an experience that has been emotionally or intellectually significant, an opportunity for strengthening bonds between friends, or a chance to broaden one’s musical horizons:

[Lending] allows others to enjoy my music, experience new types of music and allows me to share with others who have similar tastes. Additionally, it allows me to introduce lesser-known bands to my friends and allow [sic] them to enjoy the styles of music that I do.

4.8 Metadata and Extra-musical Documents Are Desired

Given the access methods described by participants, the absolute minimum metadata required to support searching

and browsing in a personal collection appears to be the artist's name, CD title, and song title. A simple way to enter these bibliographic details would greatly enhance usability of a personal music library; ideally, each piece of music would have this metadata associated with it and automatically loaded into the music digital library with the song or album itself. (for example, using a service such as Gracenote's CDDB; www.gracenote.com).

Earlier studies of music queries on the Google Answers 'ask an expert' system and on a music-focused Usenet Newsgroup suggested that a far richer set of metadata would be desired to enhance the user's interaction with a personal music collection. While this present study did not directly address the question of what metadata users would like to have available, the observations and interviews indicate that some users may desire additional metadata—for example, one participant wanted the timings for albums and songs, to allow him to keep disk changes to a minimum when using a Discman. The precise metadata desired is highly likely to vary from user to user, and so as rich a set as possible should be available, with the user able to select the fields of interest for display. As an example, consider the spreadsheet developed by one of the most avid collectors encountered in this study. This individual entered standard bibliographic details such as artist and CD title, and also details specific to his collection such as the year that he acquired each CD. Alas, even this relatively simple set of details proved too onerous to keep current, and he fell so far behind in data entry that the spreadsheet was abandoned.

The most frequently mentioned additional metadata is the lyrics of songs. Association of lyrics (or 'the words') to songs in a collection is useful in familiarizing oneself with a new acquisition: "If the CD is new I will sometimes take the insert out of the case to see if there are any lyrics printed so that I can sing along with the music."

Several participants reported using their collection to aid in musical performances, either amateur or a professional. For singers, printed lyrics are exceptionally useful, as it may be difficult to interpret the words as sung in the recording itself. For instrument players who cannot read music, the recording may be repeatedly listened to until it can be played by ear; the facility to easily repeat difficult bits until the notes are picked out would be helpful for these users. Musicians who can read music would of course benefit from having the score available together with the recording.

Enjoyment of a personal music collection may also be significantly enhanced by ready access to music-related documents, giving background or otherwise augmenting the listening experience. One participant describes a friend:

For [him], music does not begin and end with listening to the E.P., but continues onto a complete artist experience, including investigating the band on the Internet, downloading music videos and investigating their belief and social systems through thorough investigation of their official and unofficial websites.

People with such an intense interest in specific artists or genres are not uncommon; they may participate, for example, in online 'interest communities', as described in [11]. The ability to link specific songs, or groups of songs, to the information discovered online would likely be of great interest to these aficionados. Further, it may be useful to store documents about artists, albums, genres, etc. that are not directly linked to any music in the collection; information searching may be conducted prior to purchasing a piece of music, or the information gathered may indicate that a particular potential purchase would not be likely to be enjoyed by the user.

4.9. Collections are visual and tactile

Earlier work [6] describes the ways that CD cover art can be useful in searching or browsing a large CD collection—for example, when searching for a particular CD its cover can be more quickly recognized than its title, and the style of the cover art can provide clues as to a CD's genre or style. In the personal collection, these cues are particularly useful in browsing, whether looking through one's own music to find something to listen to, or when examining a friend's collection to literally 'see what's in it'.

While some participants expressed no interest in the CDs other than as a container of music ("I don't really care how it looks"), the appearance of both individual CDs and the physical collection as a whole is significant to others. In a personal collection, the CD covers may indeed be used as cover 'art':

Occasionally CD inserts with effective graphic design are used as decoration, by being U-Tacked to the wall. This allows for ease of lyric recall and adds an aesthetic element to my room.

Another participant enjoyed designing CD labels for compilation CDs that he created, as that allowed him to make the compilations more visually attractive. Still another sorted, stacked, and positioned his collection to provide an aesthetically pleasing display in his room. A collection's appearance as well as its content may be significant to 'image management' [10], how that person presents him- or herself to the world; yet another reported that,

To me, the manor [sic] at which I display my music, is almost as important as the music itself. This is part of the reason why I still retain my CD collection (as every song I have on CD, is also in MP3 format on my computer).

The sheer physicality of a CD may add to the experience of collecting music. One participant was asked why he browses through CD stores on a regular basis, given that he could simply phone the store to find out whether it has a CD that he is considering purchasing; he replied, "it's important to go and press the flesh, so to speak". Another reported that after purchasing a new CD

that he will “take it round to friends to show it off and maybe let them hold the case.” For these people, simply having an MP3 file does not give the same pleasure or the same sense of ownership, of having a collection, that the purchase of a physical CD brings.

It will be a challenge to the designers of music appliances or digital libraries, and to the music industry, to bring this sense of joy of possession to the online purchase of an MP3. One possibility for adding value is to make available other, related documents with a piece of music—for example, images, lyrics, or background information about the artist—and then to support the user in viewing or otherwise using these images through the music digital library. Personalization seems to be the key here, for example by allowing the user to easily make backgrounds or wallpaper, or to set up icons representing the piece of music.

4.10. Browsing

Browsing through a personal music collection may be extremely undirected, essentially a linear search until a piece of music suddenly strikes the individual as what s/he wants to hear at that moment: “I usually access this part of my collection by flicking through (maybe multiple times) my CD wallet looking at each individual CD trying to decide what I feel like listening to.” The end of this activity comes not when a predetermined item or type of music is located, but when a song unexpectedly attracts attention and is selected for playing.

Browsing involves scanning the CD faces if the CDs are stored in a wallet, or scanning the spines if they are stored in a tower or a stack. Spines are a sparse source of information, giving only the artist and title—although one participant reported that color might aid in recognizing a desired CD. CD faces are not always placed in the wallet so that the text is right side up, and copied CDs have only the details that the copier has thought to provide (generally very little information). Given the variety of features that may spark interest in listening to a particular piece—title, artist, genre, rhythm, and mood, to name just a few—a rich set of browsing categories is indicated.

4.11. CONCLUSIONS

A personal music digital library system will need to support mix of tasks different from those of a large, public music digital library. While significant effort is being exerted in the music information retrieval research community in developing query-by-humming interfaces, this type of access will likely be less frequently used in a personal collection. Users who interact with a set of music that they themselves have chosen for inclusion will necessarily be more familiar with its contents than with a public collection, and so will be less likely to conduct a query-by-humming search for that song that they can’t quite identify, but can’t get out of their heads. Note the

weasel words ‘less likely’; users may be able to identify the approximate location of a desired song (within a particular genre, by a given artist, or on a particular CD) but may not know the track number, title, or other identifier. In these cases, a query-by-humming interface may be useful in selecting the correct item from a set of candidates.

Commercial services for managing personal collections, such as iTunes (<http://www.apple.com/itunes>) already include a number of the facilities identified in this paper as desirable in a personal music digital library—for example, searching and browsing support over a rich set of metadata (title, artist, date, lyrics, etc.), facilities for creating playlists, and the ability to customize physical media by printing CD labels. The insights into personal music behavior coming out of this study point to areas in which the currently available facilities might be extended; for example, that the user might be able to *easily* add new metadata (and new, idiosyncratic metadata categories, such as “the person who gave this to me”, “the parties I’ve played this at”, and so forth).

Ease of use is paramount. Members of focus groups were particularly scathing about the difficulties they had encountered in using existing music management software: “...perhaps I’m just stupid, but I’m damned if I can make the thing do what I want. I mean it should be simple right?” Learning to use, and using, the system should not interfere with enjoyment of a music collection.

At present there appears to be a tension between design for small size for portability, and provision of a screen display large enough to support searching, browsing, and organization of a collection. One focus group was particularly emphatic about the need for a larger display area than currently exists on MP3 players, and for a crisp, clear display. Perhaps the current focus among manufacturers for designing ever-smaller MP3 players will lead to missed opportunities (for example, the iPod Mini bills itself as “smaller than any cellphone”); small size and portability in an information appliance should not be the primary goal in the design of an information appliance, but should be secondary as derived from user needs and requirements of function [14]. In this study, participants expressed a keen desire for many functions that would require a display of at least the size on a PDA, if not larger; no one referred to small size as important or beneficial. Not all portable music appliances need to be pocket-sized: in the past, some people took along their “ghetto blasters” to social events; in the future, people may well bring their personal music servers. Given that music collections are typically used in multiple locations, and that people are keen to enjoy a rich interaction with their collection whether using a PC or an MP3 player, it is difficult to envision a tiny-screened music appliance supporting the full set of features of a music digital library with a high degree of usability. Manufacturers who do provide small appliances meeting these challenges will be well placed, differentiating themselves in a rapidly expanding market.

The ‘Smart Playlist’ function of iTunes incorporates many of the ordering and selection features encountered in this study—inclusion by metadata values (such as genre), ordering by attributes such as newness to the collection, and so forth. Additional support for playlist maintenance could include management of the ‘thrashing’ and ‘sickness’ cycle, although it will be a challenging task to set up an appropriate interface for this feature! Given the idiosyncratic nature of genres as described by participants, in creating genre playlists it may be more natural to allow users to specify a particular song as an example of a genre and then have a playlist automatically generated based on audio similarity to the example [16], rather than asking the user to use genre metadata.

Given that a music digital library as described in this paper would be, in the words of one participant, ‘a part of your environment’, its appearance would be important to its acceptability. A further desirable extension to commercial personal music management software/systems would be more significant ability to personalize the appearance of individual songs and compilations/playlists, as well as that of the collection as a whole (in a more fundamental manner than through ‘skins’). At present it is not uncommon for a music collection to have several owners/users (for example, within a family or in a student flat); this suggests that a personal music digital library should not be strictly a single user system, but should support multiple users, each able to personalize the collection to suit their needs. A collection should also be able to be presented in a form understandable by others—to allow friends to view the collection, as part of the image that an individual presents to the world.

5. ACKNOWLEDGMENTS

We wish to thank the students for their work in gathering ethnographic data for this study; the participants for their patience and time; and the members of the New Zealand Digital Library Research Group, for their collegial support.

6. REFERENCES

- [1] Bainbridge, D., Cunningham, S. J., and Downie, J. S. “Analysis of queries to a Wizard-of-Oz MIR system: Challenging assumptions about what people really want,” *Proceedings of the 4th International Conference on Music Information Retrieval*, Baltimore, MD, USA, 2003, 221-222.
- [2] Byrd, Donald, and Crawford, Tim. “Problems of music information retrieval in the real world,” *Information Processing & Management* 38, 2002, 249-272.
- [3] Campbell, Kim. “The death of the album?” *The Christian Science Monitor*, 14 Nov 2003 edition, 14. <http://www.csmonitor.com/2003/1114/p15s01-alm.html>.
- [4] Crafts, Susan D., Cavicchi, Danial, Keil, Charles, and the Music in Daily Life Project. *My Music: Explorations of music in daily life*. Wesleyan University Press, Middletown CT, 1993.
- [5] Crawford, L. “Personal ethnography”, *Communication Monographs* 63/2, 1996, 158-170.
- [6] Cunningham, S., Reeves, N., Britland, M. “An ethnographic study of music information seeking: implications for the design of a music digital library”, *Proceedings of the ACM/IEEE Joint Conference on Digital Libraries*, 2003, 5-16.
- [7] Downie, J.S., and Cunningham, S.J. “Toward a theory of music information retrieval queries: system design implications”, *Proceedings of the 3rd International Conference on Music Information Retrieval*, 2002, 299 – 300.
- [8] Futrelle, Joe, and Downie, J. Stephen, “Interdisciplinary Communities and Research Issues in Music Information Retrieval”, *Proceedings of the 3rd International Conference on Music Information Retrieval*, 2002, 215-221.
- [9] Glaser, B., and Strauss, A. *The Discovery of Grounded Theory: Strategies for Qualitative Research*. Chicago, 1967.
- [10] Goffman, Erving. *The Presentation of Self in Everyday Life*. Anchor Press, New York, 1959.
- [11] Kibby, Marjorie D. “Home on the page: a virtual place of music community”, *Popular Music* 19/1, 2000, 91-100.
- [12] Kim, Ja-Young and Belkin, Nicholas J., “Categories of music description and search terms and phrases used by non-music experts”, *Procs. of the 3rd International Conference on Music Information Retrieval*, 2002, 209-214.
- [13] McPherson, J.R. and Bainbridge, D. “Usage of the MELDEX Digital Music Library”, *Proceedings of the International Symposium on Music Information Retrieval*, 2001, 19-20.
- [14] Norman, Don. *The Invisible Computer: why good products can fail, the personal computer is so complex, and information appliances are the solution*. The MIT Press, Cambridge MA, 1998.
- [15] Pauws, Steffen. “Effects of song familiarity, singing training and recent song exposure on the singing of melodies”, *Procs. of the 4th International Conference on Music Information Retrieval: ISMIR*, 2003, 57-64.
- [16] Pauws, Steffen, and Eggen, B. “PATS: Realization and user evaluation of an automatic playlist generator”, *Procs. of the 3rd International Conference on Music Information Retrieval*, 2002, 222-230.