Working Paper Series
ISSN 1170-487X

Toward a Theory of Music Information Retrieval Queries: System Design Implications

Sally Jo Cunningham and
J Stephen Downie

Working Paper: 05/02
June 2002

Department of Computer Science
The University of Waikato
Private Bag 3105
Hamilton, New Zealand

Toward a theory of music information retrieval queries: system design implications

J. Stephen Downie	Sally Jo Cunningham
Graduate School of Library and Information Science	Department of Computer Science
University of Illinois at Urbana-Champaign	University of Waikato
Room 103 LIS Building	Private Bag 3105
Champaign, Illinois, USA	Hamilton, New Zealand
Phone: (217) 351-5037	Phone: 64-7-838-4402
	Fax: 64-7-838-4105
Email: jdownie@uiuc.edu	Email: sallyjo@cs.waikato.ac.nz

Keywords: music query analysis, music information retrieval system design, music information needs, music information uses

Toward a theory of music information retrieval queries: system design implications

1. Introduction

Interest in the development of content-based music information retrieval (MIR) systems is growing rapidly. The MIR research community consists of a multidisciplinary amalgam of librarians, digital librarians, information scientists, computer scientists, musicologists, audio engineers, lawyers and business persons. This multidisciplinary approach has given rise to significant technological advancements in retrieval algorithms, audio interfaces and data representation schemes.

Notwithstanding these technological advancements, MIR research is currently a systems-centered research domain. For a variety of reasons—including intellectual property law, limited access to substantial, multi-genre, multi-format collections and a lack of a historical user-base—MIR research has hitherto been unable to develop and exploit data concerning the nature of real-world user needs and use of music information. Although [3] and [1] have both commented on the paucity of research involving real users of MIR systems, there have been limited attempts to garner information about and from users. The two most notable examinations of MIR system usage in the research literature are a partial analysis of the OPAC transaction logs from a music library [6] and a log analysis of text and audio queries to a testbed collection of bibliographic and MIDI files [8]. The insights gained in these studies are necessarily limited to the systems that generated the usage data—that is, the studies provided valuable hints as to the usability of features implemented in the two MIR systems, but the data could not provide insight into what additional search facilities or output types that users desire. Moreover, the testbed collection whose usage was studied in [8] is an artificial collection created to test the facilities of the MELDEX MIR software—and so it is not clear that this log analysis describes 'real' information needs.

MIR research is being hindered by its lack of user-centered theoretic or empirical groundings for its systems design decisions. It is the purpose of this paper to address this obvious need by laying out an empirically derived theoretic framework for the categorization of MIR queries through the context of the needs and uses made manifest in a collection of unsolicited and spontaneous real-world music queries. We will also demonstrate the possible design implications that emerge from our query analyses and classifications.

2. Data gathering

This study examines music information requests harvested from the rec.music.country.old-time newsgroup, an unmoderated group founded in 1995 as a forum for discussing 'old-time' music, defined in the group's charter as "music that originates in the southern Appalachian region ... but which also includes related music, for example, from Missouri, Quebec, or New England (contra)" [4]. In practice, the group also tolerates limited discussion of related genres, such as western, folk, and older country-western music.

Why analyze postings to a newsgroup? Given the dearth of query data available for analysis from formal MIR systems, we sought a source of authentic music information requests. This newsgroup provides such a source of queries, and in practice functions as a human-based MIR 'system': it is effectively an old-time music reference desk, run by volunteers keen to share their expertise in and passion for old-time music. The 'interface' to this people-powered system is flexible and usable: the users can post natural language requests, expressing their information needs in their own words and unrestricted by the artificial constraints of a search syntax. These requests contain rich contextual details and background information, permitting the researcher a greater insight into search motivations than is generally available in, for example, a transaction log from a computer-based IR system. Further, information requests to a formal IR system are often constrained by the user's pre-conception of what types of information or document formats are available—the user tailors requests (consciously or unconsciously) to what s/he thinks could be retrieved from that system. In contrast, newsgroup readers recognize that posted requests to a newsgroup could literally retrieve anything—a desired fact, an opportunity to purchase a much-desired album, a pointer to an mp3 file, and so forth. The text-based format of newsgroup postings limits this study to primarily text-based queries, and so we cannot examine in any depth the audio or symbolic techniques that people may wish to use when describing a music information need. As we will briefly touch upon in Section 3, however, the desire for multi-modal expression facilities appears even in the spare, ASCII newsgroup interface—newsgroup posters include transcriptions and links to audio files in their queries.

This particular newsgroup (rec.music.country.old-time) was selected as a source of queries for several reasons: it is an active newsgroup (by February 2002, the group's archive contained over 32,800 discussion 'threads'); the group focuses tightly on discussions of music and music events; and the discussions are relatively 'serious'—that is, the participants focus on the genre, rather than on personalities or gossip. The genre itself is also well-suited for the purposes of this study. Information needs related to old-time music can be complex to specify, as there may be no canonical or 'official' score or lyrics for a given piece, and may be difficult to satisfy, as the music itself is not widely distributed in music stores or on the Internet—either as score, lyrics or recording. Satisfying an information need related to old-time music is not a simple matter of browsing in a CD shop or running a quick search on a peer-to-peer filesharing service.

This paper analyzes 161 music queries posted during the period July 2001 – Jan 2002. During this period, approximately 760 'threads' are represented in the rec.music.country.old-time archives held by Google; each thread was manually examined, and only the 161 postings which began a thread and which contained a music-related request were retained.

For purposes of analysis, the subject line and the body of the query were extracted; the subject was retained because it frequently contains portions of the query (see example below). Initial greetings ("hello gang!") and signatures were stripped. The remaining query text tended to be lengthy, with an average of approximately 69 words per informaitn request. The longest query contained 349 words, and the shortest a mere 10 words:

> Subject: Log Roller's Waltz
> Query: Does anyone know where to find it?

Each selected posting was analyzed as a single query, regardless of the number of individual information requests that the posting contained. This strategy was adopted for pragmatic reasons; in many queries the distinction between a single request for "a song" and multiple requests for details relating to a given piece of music appears artificial. For example, is the following a single query, or three requests for three pieces of information?

> I am a preacher looking for the words/composer/artist who once sang
> "Take a Little Walk with Jesus"

The 161 postings were coded by characterization of information need (Section 3), characterization of the desired search outcome (Section 4), intended purpose or use of the search outcomes (Section 5), and additional search background details (Section 6). Note that in these codings, a single query could contain more than one categorization—for example, a single posting could ask for more than one document type to be returned as an 'answer' to the query (search outcome), or a single posting could specify more than one intended use for the desired information. The researchers cross-checked each others' codings, and consensus was reached on the categorizations.

3. Characterizing the music information need

What features or characteristics do posters to rec.music.country.old-time use to describe the types of musical information that they desire? We categorized the descriptive details contained in the postings as follows (results summarized in Table 1):

- BIBLIOGRAPHIC: the description of the information need includes music metadata; that is, the request includes information describing a musical work, rather than information or descriptive details that can be derived directly from a given composition.

- GENRE: description of the style of the desired music.

- LYRICS: the provision of some or all of "the words" to songs that include lyrics. Note that old-time music may be solely instrumental.

- LYRIC STORY: Some old-time songs tell a story, and this story may be remembered in addition to—or in the absence of—other details about the song. The description of the storyline may be brief: "a song about a lady who had tattoos of all 50 states over her body", or may be elaborate (one description was over 300 words in length).

- TEMPO: description of the speed and rhythm of the desired musical item.

- EXAMPLE: the query included representations of the music about which information was desired: for example, links to MP3 versions of songs or a transcription.

- SIMILAR WORKS: a request for works similar in style to a given example or similar to those produced by a particular performer, works following a given theme (e.g., "songs related to fishing or fish", "songs about travel (esp. about cars,) that would be fun for kids"), and additional works by a specific performer.

- AFFECT: This category includes descriptions of the 'mood' of a song or recording, or the emotional state induced by the music: for example, "mournful", "upbeat and slick", "cute", "lively", "hell-bent", "spellbinding", or "dirty".

Information need description	Percentage	Occurrences
BIBLIOGRAPHIC	75.2%	121
GENRE	9.9%	16
LYRICS	14.3%	23
LYRIC STORY	6.8%	11
TEMPO	2.5%	4
EXAMPLE	1.8%	3
SIMILAR WORKS	9.9%	16
AFFECT	7.5%	12

Table 1. Features used to describe the information need.

Discussion

BIBLIOGRAPHIC details were the most frequently provided characterizations of the music information. This category can be further broken down into the following bibliographic elements used in postings (Table 2):

- title of song

- collection title: title for an album, LP, CD, audiotape, etc.

- performer or group who created a particular recording; possibly the composer, if the original performer also wrote the song or if the song is strongly associated with one performer (who may then be perceived as its composer by its listeners).

- label: the organization which produced a particular recording.

- date: the date that a recording was produced, or the date that a song was composed.

- orchestration: an indication of the instruments and vocal parts participating in a recording.

Bibliographic detail	Percentage	Occurrences
title	35.4%	57
collection title	9.3%	15
performer	41.6%	67
label	6.8%	11
date	11.8%	19
orchestration	5.0%	8

Table 2. Further breakdown of bibliographic details describing the information need.

Specifying the bibliographic details referring to a particular recording or performance are relatively straightforward in some requests, in which the requestor is searching for a copy of a known, commercially distributed item (for example, "I'm looking for Kyle Creed's Mountain Records LP's #'s 301 and 305."). In other cases, the requestor has only fragmentary details—a performer's name, a song title. In supporting bibliographic queries through an MIR system, all of the difficulties associated with traditional bibliographic description and access also apply. While these music bibliographic problems have been considered in the research literature (for example, see [5] and [9]), this remains an open—and challenging—area of research.

References to a song or melody, rather than to a recording of that piece of music, are more problematic; for example, there may be several different titles for a song in common use, and the distinction between composer and performer may be blurred. And, of course, the poster may be uncertain or may have only a portion of the title or other bibliographic element: "'Shaving the Brush'. At least, I think that's the name""; "I think it was written...by Willie Walker and Gene Simpson. I am not sure about the names.". Dates pose exceptional difficulties for automated handling by a computer-based music retrieval system, as they are not specified by year; instead, they may be defined by decade (for example, "late 60's or early 70's"), or may be general indications of the presumed age of a song or recording: "I vaguely remember this tune from a long time ago"; "from about a million years ago".

Describing orchestration also poses difficulties for some requestors. Terminology is a problem—should the common old-time musical instrument be called a fiddle, or a violin? An instrumental effect may be known, but not the formal name of the instrument itself: "electric kind of guitar and ... a heavy metallic piano-sound as accompaniment". The description of vocal parts often includes a note as to whether the singers are male or female: "at least one man and one woman, in unison"; while performer gender is remembered, the part that they sing (alto, tenor, etc.) is not noted in a single information request.

Many of the GENRE or style descriptions included standard characterizations such as "country&western", "American folk song", and "bluegrass". Other requestors obviously required finer granularities of stylistic description than these larger categories: "almost ragtime", "oldtime tunes other than the popular fiddle tunes that cross over to bluegrass", "more folkie than bluegrass". Some information requests described categories that were immediately recognizable, but not likely to be part of a formal musical genre classification: "novelty", "comedy love song", "dirty song". Classifying musical documents into a priori genre classifications is a notoriously 'fuzzy' task even when considering well-known musical categories—and these requests include idiosyncratic genre descriptions that users would desire/expect to find represented in any MIR system that they search. One promising technique for overcoming these difficulties is to discard a priori genre classification in favor of clustering groups of musical documents by stylistic 'likeness' (see, for example, a system based on the use of a self-organizing map to cluster music files in [9]. This latter approach creates new difficulties in constructing a searching/browsing interface, however: since textual genre categories have been discarded, how is the user to efficiently specify a genre query or browse for genre categories of interest?

LYRICS provided in requests were fragmentary—generally a few memorable lines or phrases and perhaps the chorus, or portions of the chorus. Keyword searching of lyric text would provide simple and effective support for satisfying both LYRICS and LYRIC STORY requests. Note that the ability to search lyrics could also partially support some of the SIMILAR requests—those which included a description of a desired theme for a song (for example, "songs related to fishing or fish"). There is some evidence of this technique being employed in one MIR system: [8] reports that users of a searching system for MIDI files had developed the strategy of using topic-based search terms to browse the collection by theme.

Relatively few (4) queries specified a TEMPO or meter. In only one query is this defined precisely enough to be able to extract this information directly from a score: "It changes from 4/4 to 3/4 on the chorus." In other queries, the tempo is described more generally: "slow song", "lively", "old-fashion country rhythm". Again, these are imprecise terms that would be exceptionally difficult to quantify a priori—yet given an example, a listener could easily determine whether a given piece of music met his/her personal interpretation of 'slow' or 'lively'. If ever there were a situation where the use of relevance feedback techniques would be appropriate, this is it!

The EXAMPLE, SIMILAR WORK and AFFECT queries also indicate the need for sophisticated interface and query analysis components. Affordances for exemplar, similarity and affect queries are in fact sine qua non for real-world MIR systems: music is so poorly described in language that it is impossible to capture its "aboutness" or "affect". There are two central challenges to meeting user needs in this domain. First, users must be able to supply the system with the appropriate examples in whatever form they have at hand (e.g., score, MIDI, recording, etc.). However, in designing an MIR system, the level of computer literacy required to submit a query should be carefully considered—a complex interface requiring the submission of a sound file for a 'query by singing' (as required by MELDEX [7], for example) could be well beyond the abilities of many potential users. Second, there must be a mechanism for the user to specify in what ways the given examples are to be used to assess similarity or affect. For example, is it the tempo, the general "sound", the "mood or ambience" or the guitar solo that system should focus upon in making similarity judgements?

4. Characterizing the desired outcome

At this point, we examine the types of information that the requestors have specified that they would like to see as a response to their query—that is, the types of document or details that they are seeking (Table 3):

- BIBLIOGRAPHIC DETAILS: This category has been broken down further into the following sub-categories (Table 4):

    - general/unspecified : queries which include only very general requests for information to help 'identify' a song, presumably with bibliographic details such as performer/title/etc. For example: "Need info on a song called "Who Shot Willy on His Wedding Day"."; or simply "looking for a song".

    - title: queries seeking the title for a song. This type of request may be difficult to fulfil, given that an old-time song may not have an 'official' title.

    - performer: requests for the name of the artist or group who created a particular recording.

- Collection: queries seeking to identify an album, CD, or audiotape, generally by title.
- composer: requests for the name of the composer of a song.
- other: These queries include a request for the approximate date of composition for a song, the town in which several performers recorded, liner notes to an album, the instrumentation of a specific recording, and whether or not alternative tunings were used by a particular performer.

- LYRICS: postings requesting the lyrics (or more informally, "the words") to a particular song. In the majority of cases, the poster was seeking the full lyrics. In five of these queries, the requester knew most of the lyrics to the song, and required only a few lines or words.
- RECORDING: requests for a recording of a song. These requests may simply ask for "the song", or may specify the desired format. A startlingly lengthy set of formats are mentioned, including 78s, 45s, LPs, CDs, audio tapes, and a variety of digital formats (MIDI, WAV, MP3, and MP2).
- NOTATION: a request for a musical score, again in a wide variety of notations: ABC transcription, chords, guitar or banjo tabliture, sheet music, arrangement for a particular instrument (in this case, fiddle/violin), SSA vocal arrangement (soprano, second soprano, alto).
- RESOURCE: a request for the location of resources useful in locating music information—for example, websites and related newsgroups.
- BACKGROUND: a request for information about the history or origins of a particular song, performer, or genre.
- OTHER: These three queries fell outside the above categories: a request for conjectures about the object from which a novelty musical instrument had been constructed; a request for image files of the front and back covers of albums or CDs; and, intriguingly, what another poster to the newsgroup described as "the most obscure question I have ever seen posted on Usenet": "a list of the songs that were on the pavilion jukebox at Garner State Park (TX) from 1959 – 1963"[1].

Outcome type	Percentage	Occurrences
BIBLIOGRAPHIC	35.4%	57
LYRICS	29.8%	48
RECORDING	16.8%	27
NOTATION	13.0%	21
SIMILAR WORKS	9.9%	16
RESOURCE	5.0%	8
BACKGROUND	13.0%	21
OTHER	2.5%	4

Table 3. Characterizations of desired information

Discussion

These request indicate that music information seekers desire a breathtakingly wide variety of types of output, in a wide variety of representations. MIR systems exist which can provide as output BIBLIOGRAPHIC details, LYRICS, and a RECORDING or NOTATION for a song; however, few—if any—exist which encompass all these, particularly in the different recording and notational formats

[1] Astonishingly, this query elicited a response that included contact details for an individual that the respondent felt might be able to help!

requested. Not only do different people desire different types of output, but frequently a single request expresses a desire for more than one type of output: 33 requests (20.5%) asked for two or more output types (for example, both lyrics and bibliographic details). To date, extant MIR systems focus upon a limited range of available information (i.e., will return only sound files, or only bibliographic records, but not both). Given the relatively large percentage of multi-output requests, it becomes apparent that future MIR systems must be designed as information retrieval federators. For example, a system should be able to gather up relevant sound files from one source, link those files to relevant bibliographic records from another source, and link the bibliographic records in turn to possible relevant image collections, grey material collections (e.g., playlists, fan sites, etc.), and so on. Work by [2] to integrate content-based retrieval of music with the associated MARC records via Z39.50 protocol is a step in this direction.

The RESOURCE, BACKGROUND, and OTHER categories include requests for documents or details rarely held in MIR systems—historic details about a genre (BACKGROUND), pointers to what amounts to entire additional music databases (RESOURCE), and obscure music-related details (OTHER). While these types of outputs are unlikely to be incorporated into MIR systems in the near future, it is important for system designers to be cognizant of what is desired, yet currently not included.

Bibliographic detail	Percentage	Occurrences
general	9.3%	15
title	7.5%	12
performer	14.3%	23
collection	3.1%	5
composer	3.1%	5
other	3.7%	6

Table 4. Further breakdown of bibliographic details requested as a search outcome.

5. Characterizing the intended uses

Reference librarians have long recognized that understanding the intended use of desired information can be helpful in clarifying and satisfying a request. We identified the following categories of intended use in the newsgroup postings (Table 5):

- LOCATE (Default): in some sense all requests involve the location of information so any query that did not make ***readily apparent*** the intended use of the information was classified as a locate request.

- RESEARCH: requestor makes it clear that the information will be used as part of a formal or comprehensive information gathering project (e.g, "background information on fiddler Glenn Smith for a class I'm teaching") or to satisfy an personal curiosity (e.g., "also curious if any info about another band"); also included in this category are requests that ask historical questions (e.g., "origins of the song") or pose hypotheses, for example, "Does this speak to some essential difference".

- PERFORM: requestor wishes to perform the music in question, for example, "We're going to do this at a birthday party", "recording a theme album".

- COLLECTION BUILDING: the user expresses a desire to add to (e.g., "what other early country 78s fit the bill for a Holiday program", "a site where I may find more of their tunes") or start a collection of music based up some criteria, for example "similar volumes on banjo tunes".

- LISTEN: the request makes it clear that the requestor wishes to listen to the music, (e.g., "trying to hear these for years", "listen to sound files of Old-Time Music") or clearly implies the desire to hear the music (e.g., "love the way Jay Ungar plays it", "for the elderly listener").

Category	Percentage	Occurrences
Locate	49.7%	80
Research	19.3%	31
Perform	18.6%	30
Collection Building	18.0%	29
Listen	6.8%	11

Table 5. Intended uses for requested music information.

Discussion

Of the three characterizations used in this analysis, this is the most problematic. While classifying the queries in to the above categories we found ourselves making inferences and assumptions about the intended uses of the information. For example, a request for a MIDI file might lead one to assume LISTEN as the intended use; however, MIDI files are also used in both performance and research. In this case, where no other information was provided concerning use, the request was classified as LOCATE, our default category. Thus, data outlined in Table 4 is quite conservative and speaks only to what we know about intended uses and not what we might infer (albeit, perhaps correctly) about how users will be using the information provided. Notwithstanding the probable under-representation of the individual numbers for each type of use categorized, that there are non-trivial numbers of explicitly delineated instances of each category suggests the validity of the categories as analytic entities. Future theoretical and empirical investigations into music query analyses must develop more robust ways of extracting the intended uses from requestors. In the same vein, systems developers must create mechanisms through which the requestors are prompted or encouraged to more clearly explicate the intended uses of the information in question.

Two categories—RESEARCH and COLLECTION BUILDING, with 31 (18.6%) and 29 (18.0%) instances respectively—have ramifications for system development. Both of these imply (in the majority of cases) that the requestor requires more than one item and/or kind of item (e.g., scores, dates, bibliographic records, etc.). This finding underscores the importance of supporting multi-format, multi-representational MIR systems, as discussed in Section 5.

6. Additional social and contextual elements of queries

Of the 161 queries analyzed, 14 (8.7%) were information requests made by the poster on behalf of another person. The reasons given for posting for another person varied. Some of the queries mention acquiring the musical information for an elderly friend or relative, who presumably had no internet access or was not computer literate enough to post the query personally. One query explicitly stated that the query was on behalf of a friend who, while a skilled craftsman, was reluctant to gain computing skills: "He makes his own dobros and mandys but don't want to do this nerdy web stuff. HE MAY HAVE A POINT."

Other requestors appear to be searching for the music or musical information as a favor or gift for the recipient: "If anyone would help me I would sure make an old man happy"; "It's for a friend's recording project". These personal interjections surely must increase the likelihood of a response to the query—one would have to be hard-hearted indeed to resist the opportunity to add to an old man's happiness! On the other hand, a selfish or unethical presumptive reason for seeking information may decrease the possibility of a response; the request for cover images for 17 recent-release music CDs was condemned ("Get a grip. Like anyone here would help you pirate music CD's.").

Twenty queries (approximately 12%) contained indications of specific resources that the requestor had examined in an attempt to satisfy the information need. Three requestors had turned to other people for help (an instructor, an "informant", and other people "in my area"), one poster had examined paper resources, one poster had previously consulted newsgroups or other online discussion groups ("I occasionally post this inquiry in one place or another, and have yeat [sic] to get any reply"), and the remaining 15 queries described conducting WWW searches or examining other WWW or Internet resources ("websites" in general, a WWW auction site, Napster, online stores, and music-related websites). This type of detail could affect the probability of a useful response by underscoring the importance of the query from the perspective of the requestor through a discussion of the thoroughness of the search, or through a description of the time invested in the search ("I have been looking for the chords to this terrific Red Clay Ramblers song for years"). Further, the potential respondent is saved the effort of consulting a resource that the requestor has previously examined and knows is not useful.

The associative or environmental context of the desired music-related information was mentioned in 18% (30) of the Usenet postings analzyed. By this, we mean that the posting mentioned a social connection that brought the music to the requestor's attention, or an emotional/social memory associated with the music, or that the music is associated with an event other than simply listening to a CD or LP ("I've heard it at a couple of jam sessions", "ex-friend of mine ... used to sing a song when bowling", "a mandolin class").

In some cases, this associative context may be helpful in satisfying the information need; for example, the detail that a song was heard by the requestor "Last week at the Fraley festival" is potentially useful to responders, as the responders may have also attended the festival, may have access to the festival programme, or may use background knowledge about the type of music included in that festival to identify the song in question. The mention that a song featured in a movie or appears on a website of interest to the newsgroup's readers is also likely to provide significant clues to a responder. Highly emotional associations may serve the purpose of building an empathetic link between the requestor and potential respondents; certainly the emotional motivation prompting some requests would encourage a reply ("This song was one of my Dads favorite songs and my Mom used to play the guitar and sing it for Dad. They bot [sic] passed away in the last 4 years; and I would love to have the words to this."). In other cases, the associative context appears less useful in answering the requestor's query, and may simply be part of a stream-of-consciousness style of expressing the information need ("I heard this song on the radio a couple years back and have not heard it since").

7. Summary and future research

This paper analyzes a set of music information queries across several facets: we examine the types of details presented to describe the information need, the types of information that are desired to satisfy the information need, the explicitly described intended uses for the information, and the social/environmental/associative contexts mentioned in the request. This study suggests that findings from studies of music information requests can be used to inform the development of effective, usable MIR system interfaces, as well as to indicate the types of documents or document representations that a given group of users would require to be included in an MIR system for it to be capable of satisfying their information needs.

However, these queries are drawn from a single source, and describe music information needs based around a single genre. Further studies from a variety of sources of music information requests, across a variety of musical genres drawn from diverse cultures, are required to validate to validate the categories presented in this paper and to discern broad-based statistical patterns within the categories. The hundreds of Usenet music groups are a rich source of music query data, as are the multitudinous email discussion lists and websites dedicated to particular types of music or groups of music appreciators. It is also important to move beyond the limitations of text-based query description, to capture the multi-modal richness of real-world, spontaneous "music" inputs—for example, to record and analyze queries including hummed and sung snippets of songs. These latter types of queries are more difficult to obtain; we plan, for example, to negotiate with producers of a popular radio programme to gain access to the musical queries phoned in to their on-air experts.

References

1. Byrd, D., and Crawford, T. (2002) "Problems of music information retrieval in the real world". *Information Processing & Management* 38(2), 249-272.

2. Dovey, M. (2001) "Adding content-based searching to a traditional music library catalogue server". In E.A. Fox & C.L. Borgman (Chairs), Proceedings of the 1st ACM/IEEE Joint Conference on Digital Libraries (Roanoke, Virginia, USA; June 24-28, 2001), ACM, 249-250.

3. Downie, J.S. (in press) "Music information retrieval". To appear in *Annual Review of Information Science and Technology 37*. Medford, NJ: Learned Information.

4. FAQ of rec.music.country.old-time, http://home.earthlink.net/~stevesag/otfaq.html, viewed Feb 21 2002

5. Howard, J., and Schlichte, J. (1988). "Repertoire international des sources musicales (RISM)". In W.B. Hewlett & E. Selfridge-Field (Eds.), *Directory of computer assisted research in musicology 1988*, 11-24. Meno Park: Center for Computer Assisted Research in the Humanities.

6. Itoh, M. (2000) "Subject Search for Music: Quantitative Analysis of Access Point Selection". In Proceedings of the 1st Annual International Symposium on Music Information Retrieval. Amherst, MA: Center for Intelligent Information Retrieval. Available: http://ciir.cs.umass.edu/music2000/..

7. McNab R.J., Smith L.A., Witten I.H. and Henderson C.L. (2000) "Tune retrieval in the multimedia library". In *Multimedia-Tools and Applications 10*,113-132. Kluwer Academic Publishers.

8. McPherson, J.R., and Bainbridge, D. (2001) "Usage of the MELDEX digital music library", Proceedings of the Second Annual International Symposium on Music Information Retrieval (Bloomington, IN, USA; October 15-17, 2001), pp. 19-20.

9. Rauber, A., and Frühwirth, M. (2001) "Automatically Analyzing and Organizing Music Archives". In *Proceedings of the 5. European Conference on Research and Advanced Technology for Digital Libraries (ECDL 2001)*, Sept. 4-8 2001, Darmstadt, Germany, Springer Lecture Notes in Computer Science, Springer.

10. Temperley, N. (1993) "The problem of definitive identification in the indexing of hymn tunes". In R.D. Green (Ed.), *Foundations of music bibliography*, 227-239. Binghampton, NY: Haworth Press.