Digital Tools for Musicology

A Methods Network Working Paper

This paper is a summary of, and a reflection on, some of the current commentary about the use of ICT methods in the field of musicology and (in common with the other working papers in this series) relies significantly on the contributions and ideas that have been expressed by participants at a number of seminar and workshop events funded by the AHRC ICT Methods Network. The principle purpose is to examine a selection of recent and current technical approaches to musicology with a view to promoting and encouraging the use of ICT techniques for research within the discipline. The range and complexity of computational approaches to musicology is a daunting prospect for researchers who have - for whatever reason - been unable to progress their academic interests using digital tools, but it should be clear from the following account that there is an ongoing, significant and sophisticated cross-disciplinary effort to open out areas of research within the subject. It is also clear that the ultimate objective of many of these initiatives is to more clearly define how technology can serve the research agendas of the music studies community, and how advances in technology will impact upon those objectives, in terms of efficacy, scale and achievability.

Musicology and ICT Methods

In all arts and humanities disciplines there are areas of enquiry that more or less lend themselves to computational methods of research and the field of music is no exception. For the purposes of the 2001 Research Assessment Exercise (RAE), the fields encompassed by the discipline of 'music' include:

> 'Composition and performance (including classical, commercial, and popular); history and criticism of music; ethnomusicology; theory and analysis, including empirical approaches; technology and computer applications.' [1]

This is one way of breaking down the complex relations between sub-disciplines but one could equally talk in terms of the different domains of music as being either acoustic (physical), auditory (perceived) or graphemic (notated).[2] Definitions can be problematized further by consideration at any length of concepts such as 'composition' and 'performance'. The opportunities afforded by software to analyse and then re-engineer existing music blur traditional definitions of the first term, whilst the latter has to support an enormous diversity of potential investigative approaches including: physiology, psychology, acoustics, material studies, cognition, perception etc…

On the face of it, the patterns and relationships that are discernible between the relative pitch, intervals and timbre of musical sounds would seem to lend themselves well to formal and quantitative analysis methods but this simple formulation underestimates the complexity of the language and expressiveness of music, and perhaps overestimates the precision and flexibility of digital analysis and retrieval techniques. This conception of what computers can do in the realm of musicology is also far too limiting. For the purposes of this paper, the term 'musicology' is understood quite broadly and will encompass digital techniques that impact upon historical, visual and cognitive approaches to the subject, in addition to the type of analysis that might generically be termed 'digital signal processing'. (Performance related issues will be excluded by nature of their alignment with practice-based research, a topic that will be covered in another paper).

User-led Approaches

[1] RAE 2001 results at http://www.hero.ac.uk
[2] Cited by Wiggins (2006) from Babbitt (1965)

The emergent popularity of a field of research that is devoted to the analysis of music as a series of acoustic sounds can be largely attributed to the shift from analogue to digital techniques over the last quarter of a century. For the average researcher without recourse to specialist equipment and techniques, the range of analysis tasks that could be accomplished with analogue tape were severely limited. Practically speaking, the user could either listen to the recording or cut it up and splice it. Undoubtedly this is an inaccurate summary in relation to groups engaged in pre-digital era commercial, industrial or scientific sound research, but in the context of academics working in or with university music departments, the restrictions will have been significant. As a result, it is only relatively recently that the history of musicological research has begun to diverge from a focus on the manuscript or ‘hard copy’ version of the musical notation, resulting in a body of research that previously drew extensively on palaeographical and philological methods.

Inevitably, the emergence of increasingly powerful desktop applications has enabled the musicological community to engage with a wide range of technical approaches and in common with most other disciplines, progress towards standardisation of resources and a general acceptance of key tools is broadly still to be realised (with the exception perhaps of XML, which functions as a fundamental encoding system across a wide range of activities). What has undoubtedly made an enormous impact on the field however is the ubiquitous popularity of the MP3 format and the associated hardware and software that was produced in response to the appearance of a compression format that made web transfers of audio files a realistic proposition. As a bottom-up approach to archive building and resource creation, it is clear that software such as *iTunes, MusicMatch* and *Napster* and sound editing and recording packages such as *Audacity, Pro Tools, Cubase, Wavelab* and *Garageband*, will continue to have an effect on the way that a broad community of users will engage with audio information, from the professional sound engineer using *Pro Tools* in the context of well equipped studio down to the individual researcher building up an audio study library for personal use on his iPod. At the novice level, as an introduction to the principle that digital sound files can be visualised and particular passages identified, selected, edited, looped and adjusted, the waveform view (based on signal output or volume) that all sound editing software packages feature, can be a very instructive and useful overview of the entire audio file (see fig.1).

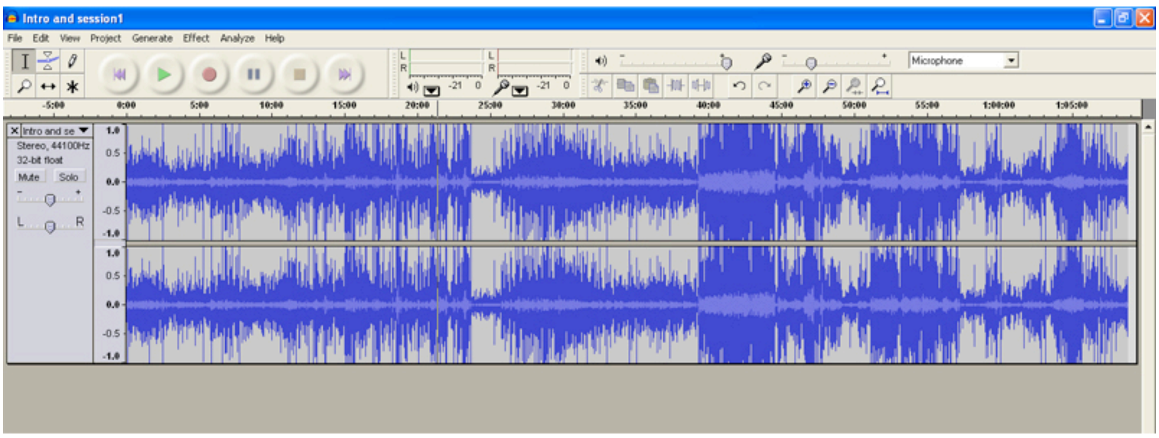

Fig.1 An audio (mp3) file as represented in *Audacity*

Another commonplace technology that has had a significant impact on the discipline, and has been part of the electronic music landscape since the 1980’s, is MIDI (Music Instrument Digital Interface). It is, inherently, an abstracted and incomplete form of musical representation but it continues to play a valuable role as an enabling technology for a very wide range of recording processes and also acts as the data source for simple ‘piano roll’ visualisations of a musical score.

This simplified benchmark representation of the score, involving data about the onset, the intensity (i.e. the velocity and therefore the volume of the keystroke), the pitch and the deterioration of notes, can be useful when analysing similarity and variation between versions of the same piece of music. As a protocol for the transmission of data, the function of MIDI can in fact be applied wherever timed electronic functions need to be applied to equipment that supports the MIDI interface (e.g. show control, theatre lighting and special effects). As later examples will demonstrate, musicology software often seems to take the form of toolkits and suites of tools that are brought together to provide aggregated functionality and in many cases, MIDI is a component part of that methodological approach, much like XML which provides developers with another (more sophisticated) representation and annotation structure.

Internet and Database Resources

At a recent Methods Network workshop[3], Michael Casey remarked that the ideal web resource would learn to serve the user's requirements rather than force the user to learn the requirements of using the resource. The potential for building 'intelligence' into systems is currently a focus of much research across disciplines and can be seen in various applications of 'probabilistic' matching and tagging systems. In the field of linguistics, a team at Lancaster University are developing a tagging process that will apply probable variant spelling tags to a large dataset containing historical word forms, based on a limited sample dataset that has been manually marked up to provide the blueprint for a subsequent and more substantial automated process.[4] In the field of content-based image retrieval (CBIR), recent progress has been made by applying feedback to systems that attempt to supply similarity matches for a query image. The retrieval of meaningful records based on the formal properties of an image alone (rather than on any associated textual metadata) becomes increasingly more difficult as the system requires more semantic understanding of the objects depicted and the input of relevance data based on sample searching gives CBIR systems an opportunity to then apply probabilistic matching to the rest of the dataset.[5]

Closer to the realm of musicology (in that there are significant holdings of music material there), the British Library uses a method that relates to these intelligent retrieval techniques. The website logs all the searches that users try to carry out, whether the attempts to retrieve information are successful or not, thereby building up a useful picture of how users intuitively prefer to try and find what they are looking for. By analysing these terms and the search results accrued by using them, it will then be possible to reformulate system ontologies to accommodate terms and ideas that were not originally perceived to be appropriate but which have subsequently proved to mean more to the general public than the system developers anticipated.

During the recent Goldsmiths workshop[6], Matthew Dovey offered further opinions about how some of the current shortcomings of the British Library catalogue might be addressed and these related to remarks made about the enormous amount of retrospective material contained in the library catalogues that had insufficient indexed metadata associated with the records. His proposal took the form of 'moving the code rather than the data' and involved writing new algorithms which would go in and search available but unindexed data to answer specific and important research questions, using grid and agent technology. This bespoke approach to finding data is clearly an unusual model for arts and humanities research but is more orthodox in the context of science and e-science disciplines, where it is more acceptable to allow substantial research queries to take hours or days to return results. Stuart Jeffrey from the Archaeology Data Service recently reported that it was routine for that organisation to run a batch job overnight that ran all possible queries that a user might make on a click and browse system containing over a million detailed archaeological records.[7] This sort of data processing approach provides arts and humanities researchers with an alternative way of looking at information management and may be of particular interest to musicologists due to the size and complexity of audio file formats and the processing overhead that is associated with carrying out even moderately complex analysis jobs.

---

[3] The Future of Information Technology in Music Research and Practice, a Methods Network Workshop, Goldsmith's College, 8/9/2006 http://www.methodsnetwork.ac.uk/activities/wsp6.html
[4] The VARD (Variant Detector) tool. See http://ucrel.lancs.ac.uk/events/htm06/
[5] Eidenberger (2004)
[6] Workshop 8/9/2006 op. cit.
[7] Jeffrey (2006)

As might be expected, the number of resources for musicologists on the web is prodigious and it is outside of the scope of this paper to comprehensively summarise them. In broad terms though, despite all of the musical resources held in significant collections such as the National Sound Archive, the OCLC Music Library and the SONIC Library of congress catalogue, it is apparent from the frustrations voiced by academics working in the field that there is no accessible and coordinated repository of audio source material which will act in the same way as the large reference corpora that exist in the field of textual studies and linguistics.[8] Tim Crawford recently remarked in the context of carrying out analysis using corpus search and retrieval techniques, that even if it were possible to analyze all of Mozart's 41 symphonies, there is no reference corpus available against which one could compare and contrast the results.[9]

One notable web resource that is enormously important to the study of an earlier period of music is the Digital Archive of Medieval Music (DIAMM),[10] an AHRC and Mellon funded initiative that has amassed a collection of around 7000 high resolution images of polyphonic sheet music. It is significant that this resource is image-based rather than audio-based, but in terms of the standards of image capture that are employed and the value added features available on the website, it represents a benchmark implementation of how to represent historical music-related material on the Web. The highly effective image zoom function, the user annotation area, the contextual help feature, the amount and quality of image metadata associated with the records and the links to other resources such as RISM (Répertoire International des Sources Musicales), all combine to add functionality and usability to the system. The embedded expertise in digital restoration techniques that is apparent in the presentation of some of the more damaged documents; and the explicit demonstration of those techniques in online tutorial pages is also extremely useful. (A Methods Network workshop organised by Julia Craig-McFeely on 'Digital Restoration for Damaged Documents' was a by-product of DIAMM's expertise in this area and has resulted in a workbook detailing relevant techniques in Adobe Photoshop).[11]

Annotation and Representation Tools

Sheet music notation as constructed out of clefs, staves, bars and notes is a highly effective distillation of the idea of a musical work but is seriously limited in terms of what it can faithfully represent about the expressive qualities of a performance, e.g. timbre, changes in tempo and very specific interpretations to do with the onset and decay of notes. One of the key areas for musicologists interested in uses of ICT is the representation of music and what decisions need to be made to capture the essence of a performance or a recording to a scholarly and precise level of detail. This matter of degree is a core issue and encompasses debate about the amount of abstraction and reduction that is necessary to adequately represent the music without losing any of the key elements that make the relevant piece of music interesting, irrespective of whether that interest lies in the context of the historical, performance-related, technical or cultural domain. These decisions are central to the strategic policy of archives and repositories and are obviously bound up with universal and ongoing discussions about standards more generally, not just for musical resources but wherever digital material is being acquired and archived.

Potentially, the most useful long-term approach to the representation of musical resources in archives would be (unsurprisingly) to include as much data as possible in relation to both the recorded signal (acoustic or electronic) and the available complementary data that provides the context for the production of that signal. The first of these involves two main considerations: the sample rate at which the signal is recorded; and the compression ratio used by the format in which the file is saved. The compressed file acts as an index and surrogate for the original performance and once that performance is finished, the resultant

[8] E.g. the British National Corpus. http://www.natcorp.ox.ac.uk/
[9] Workshop 8/9/2006 op. cit.
[10] http://www.diamm.ac.uk/
[11] http://www.methodsnetwork.ac.uk/activities/act5workbook.html

recording is the best approximation that researchers have to reality; and it is therefore critical to many areas of scholarship that as little of the original information is discarded as possible. Regarding complementary data, the same principle applies but might be formulated in a slightly different way. Having the clearest audio source possible will undoubtedly be of service to subsequent generations of researchers; it is less clear however which questions future researchers will want to ask of the metadata and considerations of what to exclude will always need to be taken in light of emerging patterns of research and technological advances.

In practice, capturing information at the comprehensive level of detail that the comments above endorse is an expensive and substantial commitment, but there are nonetheless tools available which attempt to represent music in a scholarly way, some of which are based on XML approaches. The literature on these applications is extensive but a good summary of some of the packages that have been developed can be found on a website hosted by OASIS (the Organisation for the Advancement of Structured Information Pages).[12] Out of all of the packages that have been developed, it might be instructive to look more closely at two examples, MusicXML and MEI (the Music Encoding Initiative).

MusicXML is commercially developed software[13] and its purpose is to allow the interchange of musical data. The website introduces the system as follows:

> MusicXML is a universal translator for common Western musical notation from the 17th century onwards. It is designed as an interchange format for notation, analysis, retrieval, and performance applications. The MusicXML format is open for use by anyone under a royalty-free license, and is supported by over 60 applications.[14]

The key function of this system, based on a set of 13 document type definitions (DTD's), is to allow users a way of producing printed scores from intricately encoded XML documents encompassing a very wide menu of elements, attributes and entities that attempt to precisely annotate both the acoustic and descriptive components of a piece of music. [15] As stated above, the encoding system is interoperable with a wide range of other applications, including the market leading commercial notation packages *Sibelius* and *Finale* and also the widely used academic tools, *MuseData* and *Humdrum*.

MEI is the copyright of Perry Roland and University of Virginia and also relies on a very detailed set of DTD's to facilitate precise description of musical works.

> It is designed to be comprehensive, that is, it provides ways to encode data from all the separate domains, i.e. logical, visual, gestural (performance), and analytical, commonly associated with music.[16]

The design and development of MEI was influenced by the principles that guided the creation of the Text Encoding Initiative (TEI). Both approaches do not assume that the object of the encoding is by definition a specific type of entity (i.e. a particular type of written work or a particular facet of the musical process), but there is an underlying assumption that the original entity can be expressed in a *written* form. The scope and flexibility within MEI means that the music community is free to encode a wide range of materials using the system including critical editions and collections of works containing extensive amounts of music-related text; material which would usually lend itself to TEI encoding methods.

Where possible, MEI uses familiar terminology for elements and attributes, allowing the marked-up text to be both machine and humanly readable. This accessibility to users without advanced technical software

[12] http://xml.coverpages.org/xmlMusic.html
[13] http://www.musicxml.org/xml.html
[14] http://www.musicxml.org/xml.html#Files
[15] http://www.musicxml.org/xml/musicxml-index.html
[16] http://www.lib.virginia.edu/digital/resndev/mei/

knowledge is a critical factor for the widespread use of XML (and its related techniques) across all disciplines, but it is also extensible enough to enable MEI to cover imprecise or historical musical phenomena, such as the encoding of variant readings and support for mensural music and non-aligning bar lines. In combination with the fact that MEI files can be derived from MusicXML data and can be transformed into MIDI format (and can be also be converted to postscript and pdf renditions using XSLT), there are compelling reasons for putting further resources and effort into developing this system to accommodate even more sophisticated descriptions of musical data, appropriate for the purposes of scholarship

Despite the persistence of XML developers in their attempts to provide a de facto (or actual) standard for musical representation, aspects of that representational process continue to elude satisfactory definition. The principle problem for some researchers is that musical knowledge might be characterised as multiply hierarchical and may require the sort of overlapping and cross-category descriptions that XML is not designed to accommodate, based as it is on tree-like structures. Geraint Wiggins has stated that using non-standard features of XML to try and cope with this complexity (e.g. milestones) should be considered bad engineering. He argues that this term would be equally applicable to a situation where one tried to implement whatever mechanism one had devised to compensate for XML's shortcomings *on top of* an XML structure, in order to try and adequately represent the original data.[17] In proposing an alternative model for musical description, which he broadly refers to as 'knowledge representation', he describes an approach that would seek to increase the scope for both expressive completeness and structural generality, two ideas that are normally more or less mutually inverse. His contention is that the minimal expressive power required to represent the multiple hierarchies of music is the directed graph (see fig.2) and that this should be used in conjunction with XML which would facilitate data interchange of syntactic descriptions of the graphs.

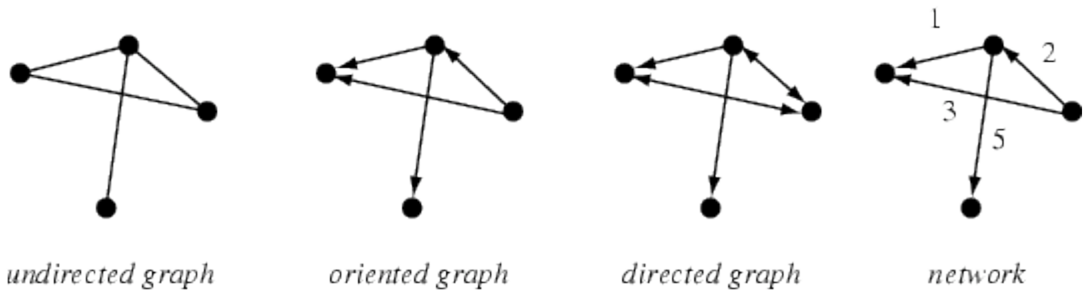

Fig. 2 (taken from: http://mathworld.wolfram.com/Graph.html)

It must be allowed that the type of 'knowledge representation' language that Wiggins refers to raises more questions than it answers but it is one of a number of attempts to lay out theoretical pre-requisites for a comprehensive representation system for music that accommodates both engineering and philosophical issues.

An alternative method of representation is to display digital images of printed musical scores and to enhance their usability with web and database techniques, as featured in the Online Chopin Variorum Project.[18] Using CSS (cascading style sheets) and javascript, the website offers an unusual level of user interactivity in that bar-length details (generated on-the-fly) from a page of music can be dragged around the screen and superimposed over each other for comparison purposes. In the pilot phase of this project, research was also carried out into providing a third-party online annotation system and some progress was made in scoping out the problem. In the second phase of this project (end April 2008) further investigation into this area is scheduled. The same team[19] is also working on a project to create a substantial archive of

[17] Wiggins, G.(2006)
[18] http://www.ocve.org.uk/
[19] Royal Holloway, University of London and CCH, King's College London

around 5000 digital images representing Chopin's first editions, which will interlink with information taken from John Rink and Christophe Grabowski's *Annotated Catalogue of Chopin's First Editions*.[20]

Music Information Retrieval (MIR)

Consideration of MIR tools separately from reflections on methods of musical representation is a rather arbitrary distinction, closely entwined as the two activities are. Nonetheless, MIR is still referred to as an 'emerging discipline',[21] and as such, conveys the sense that the use of the tools and methods associated with this activity might sensibly and usefully form a sub-discipline in its own right. In practice, some of the broader issues that demand attention from MIR specialists replicate the questions already posed to do with how music can be adequately represented. Alain Bonardi in a paper at ISMIR 2000[22] enumerated a list of musical entities that would need to be considered by the contemporary musicologist which echo the domains laid out by Babbitt in 1965.[23] However, where MIR practitioners are seeking ways of analysing pre-existing data that was captured without reference to the type of concerns that are currently being voiced within the MIR community, retrieval research can be considered a discreet practice and one that requires complex and sophisticated toolsets, many of which are daunting to the non-specialist and reflective of the historical involvement of computing science in the development of musical analysis tools.

The application that is widely regarded as the standard academic tool for music analysis is *Humdrum*, developed by David Huron well over ten years ago and originally based on a UNIX model of application development and command-line driven functionality. The FAQ section on the main Humdrum website makes it clear who the target user group for the application is:

> Humdrum is rooted in the UNIX "software tools" design philosophy. That is, each tool in the toolkit carries out a simple operation. However, by interconnecting the tools, the capacity for music processing is legion. In essence, assembling Humdrum command lines amounts to a form of computer programming. Learning Humdrum is comparable in complexity to learning pascal, perl, or kornshell programming.[24]

That complexity is undoubtedly a disincentive for many musicologists and it is clear that Humdrum remains an underused (if highly regarded) tool.

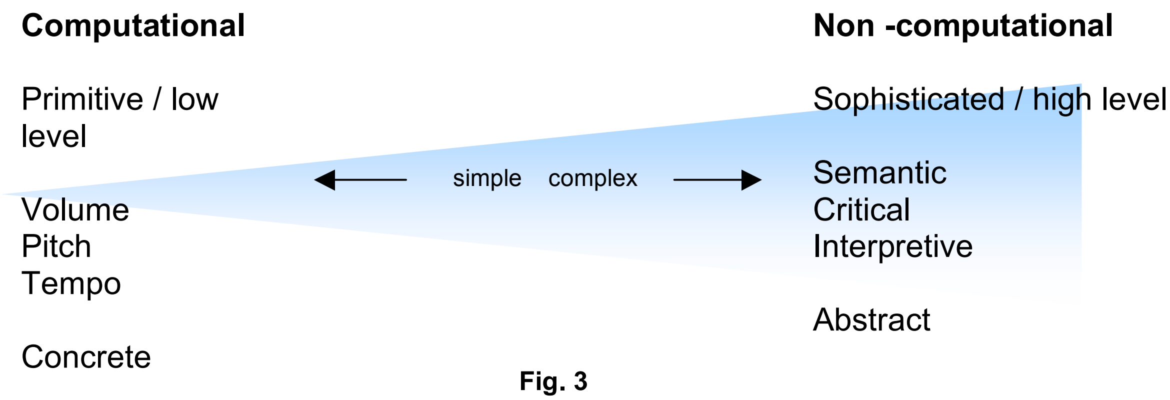

Fig. 3

---

[20] http://www.cambridge.org/catalogue/catalogue.asp?isbn=0521819172
[21] Fingerhut, M., (2006)
[22] Bonardi, A., (2000)
[23] Bonardi lists these as 'graphical', 'sound', 'symbolic' and 'tool' representations; all of which can be mapped onto what Babbitt refers to as 'the three usual representations of musical experience', the 'acoustic', 'auditory' and 'graphemic'. Bonardi, A., (2000) and Babbitt (1965)
[24] http://www.lib.virginia.edu/dmmc/Music/Humdrum/humdrumFAQ.html#hum38

The power and the flexibility of the application comes at the usual cost of intuitive usability and this kind of binary opposition is echoed in any kind of visualisation of the matrix of difficulty that one faces in processing different kinds of signals, a paradigm that is central to the problems of MIR and reflected in a range of other disciplines (see fig. 3).

Low level musical features such as loudness, pitch and brightness are relatively easy to quantitatively analyze but aggregating that analysis for the purposes of determining entities that might describe timbre properties or melodic passages quickly becomes difficult. One of the 'holy grails' of musicology is to devise an effective method of automatically transcribing music from a complex (polyphonic) acoustic source straight onto a page in musical notation. Despite a great deal of work in this area, current effectiveness begins to falter at the level of a solo piano and the prospect of capturing simultaneous multiple instrumentation is still an open problem. High level features (in common with so called 'level 3' features in the realm of image retrieval[25]) involve human perception and interpretation and as such are generally outside the scope of automated computational processes. This doesn't of course preclude methods of hybrid analysis that combine automated low-level analysis functions alongside the capacity to record manual interventions and annotations at higher perceptual levels, and applications do exist that include this kind of flexibility.

A tool that has been in development since around 1996 and has the status of an ISO standard is the MPEG-7 package, a comprehensive multi-level multimedia content description framework that uses metadata and is expressible in XML. As a flexible apparatus for carrying out the kind of research in question, it would appear to be ideal – and not only for musicology. As Adam Lindsay states, 'if someone were to invent a framework serving the arts and humanities research community for its metadata needs, it would resemble MPEG-7, at least conceptually'.[26] Acceptance and wide usage of the standard has been inconsistent however, undoubtedly due to the complexity of applying its mechanisms to actual applications. Some development has been carried out however and one example is the MPEG-7 Audio Encoder which has an associated graphical user interface developed by Holger Crysandt.[27] This tool provides users with a simple way of producing a large amount of very complex XML encoded information and this alone is instructive about the scale of the problem of processing audio information. Even when restricting the analysis to low level features[28], the processing of a 5.7 MB mp3 file results in the generation of an XML text file of 1200 pages (11MB) as viewed in a word processor application!

The development of the MPEG-7 standard was influenced by advances in 'query by image content' and 'query by humming'[29] retrieval models that were being developed in the mid-1990's and the remit of the standard grew as more applications of the technology were encompassed by the development group. The visual part of the standard concentrates on signal processing and very compact representations whilst the Multimedia Descriptions Schemes subgroup (MDS) opted for a very rich and complex set of description structures. MPEG-7 Audio went for a middle path which included options for using both high-level and generic, signal-processing-inspired descriptors.

Michael Casey has described the development of another tool based on the MPEG-7 standard, called MPEG-7 Audio Codec (MP7AC), which is an attempt to tackle the difficulty of analysing polyphonic music.[30] The objectives in the tests described are to identify similar musical sequences from two contrasting collections of classical piano works and pop songs. The features used in the analysis roughly correspond to measures of timbre and harmony. The retrieval of relevant information across any collection greater than

[25] Eakins (2000)
[26] Lindsay (2006)
[27] http://www.ient.rwth-aachen.de/team/crysandt/mpeg7audioenc/MPEG7AudioEncApp.jnlp
[28] The low level analysis features in this instance are: audio power type, audio signature type, audio spectrum centroid type, audio spectrum envelope type, audio spectrum flatness type, audio spectrum spread type
[29] A simple 'query by humming' interface is available at: http://www.musipedia.org/query_by_humming.0.html
[30] Casey (2006)

a few works is, however, computationally very onerous and another tool was developed, MPEG-7 Audio Retrieval (MP7AR), to enable searching to take place using locality sensitive hashing (LSH), a process that Casey describes as statistically dividing a feature space into regions of similarity and allowing searching to occur only within those regions. As with most musical analysis, similarity searching is open to varying degrees of specificity depending on the nature of the research, from audio fingerprinting where the concept of similarity has a very narrow definition to the concept of genre labelling where categorical boundaries are fuzzy and will involve the use of unspecific retrieval attributes.

Both MPEG-7 and Humdrum are toolsets rather than specific resources and this model is pervasive in musicology tools development[31]. The Music Lab 2 (ML) project developed at IRCAM (L'Institut de Recherche et Coordination Acoustique/Musique) is a suite of software applications to support the study and teaching of music. ML-Annotation builds on the research done at the centre into a wide range of audio related challenges and offers users a visualisation and semantic and syntactic annotation tool for use on multiple views of works which might include: musical scores, composer's sketches, recordings, other scholarly annotations, related textual material, and so on. It is designed to work as a standalone desktop resource or across distributed online libraries and works in conjunction with other tools in the suite, ML-Maquette and ML-Audio, all of which together provide representation, analysis, retrieval, playback and even compositional functionality.

Analysis by Composition

The logical extension to research techniques that progress through stages of capture, representation, retrieval and analysis is to then test that analysis by playing back and comparing datasets with original and comparative source materials. Bonardi states, 'The musicologist is at the same time a listener and a composer, since analyzing a piece of music leads to "rewriting" it'.[32] The range of software on offer to composers begins to stray out of the scope of this paper but like many areas of musicology, boundaries are difficult - and in some senses unnecessary - to define too closely. Software such as Max/MSP and SuperCollider are powerful and sophisticated composition tools based on the object-oriented programming environment and will inevitably involve practitioners in analysing the components and structure of both the pieces that they create and those of others. Another framework for composition, also developed by IRCAM, is OpenMusic 5 [33], a complete programming language that underpins the ML-Maquette application, another example of the layered and blended approach to application and function building within the music software community.

The OpenMusic environment allows for graphical representations of very complex entity relationships showing musical objects, probabilistic rules and operational elements as blocks and wireframe connections. The visual language is built on the Common LISP Object System and users are provided with a number of basic 'classes' and generic functions which represent musical structures such as notes, chords, sounds, break-point functions etc. The user then augments those 'classes' with their own defined objects and sets inheritance relationships between all the entities. Combining programming functionality with an in-built music notation editor (and instant playback features) offers the user a truly dynamic research environment where creative analysis blurs into composition in exactly the way that Bonardi describes.

Taking the theme of compositional analysis a step further, the field of algorithmic composition is a burgeoning area of activity and encapsulates many approaches to music creation. The Live Algorithms for Music[34] (LAM) group has links and is interested in multi-disciplinary collaborations that encompass work in

[31] One example of an application built on the Humdrum toolset is 'Themefinder', an online tool for discovering musical themes from a given dataset.
http://www.themefinder.org/
[32] Bonardi (2000)
[33] Bresson et al (2005)
[34] http://doc.gold.ac.uk/~mas01mc/LAM/

areas such as neural networks, evolutionary algorithms and swarm theory. One specific example that indicates the nature of this type of research is composition using cellular automata (CA) techniques and Java programming.[35] Based originally on research that focused on building visual depictions, audio CA uses the same technique of generating a new entity (or cell) based on the properties of the previous cell that was created and the status of the adjacent cells that it finds itself next to. With visual CA models, time increases as you go down the page and this largely corresponds to representations of music, so the transposition from one mode to the other is compelling and can produce complex and interesting results.

Another interesting field of research that straddles compositional and analytical approaches is the use of spectrograms (also known as sonograms) to visualise audio material (see fig. 4). These graphical views represent the frequencies of audio material and can be generated 'in reverse', starting with an image in order to produce audio output that corresponds to the elements on the x and y axes where x equals time and y equals frequency.

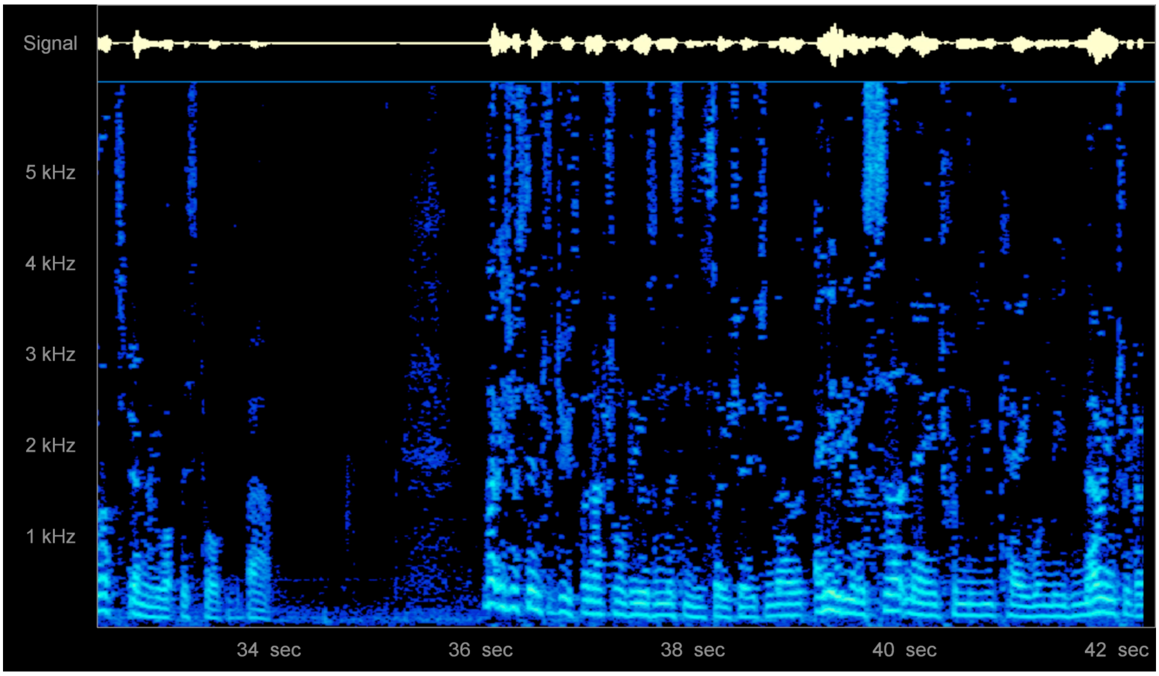

Fig. 4 (screenshot taken from Spectrogram 14)[36]

This tool is a featured component in various analysis software suites but has also been used by musicians such as Aphex Twin to 'embed' images within the sequence of audio tracks presented on commercially produced compact discs.[37]

Conclusion

In tackling some of the issues that are currently being discussed in the realm of musicology, many other potential avenues of enquiry are opened up and opportunities for cross-disciplinary engagement with the subject become abundant. The role of the musical critical edition is one such area that clearly invites collaborations between information scientists, historians, imaging experts, textual scholars and musicologists. The approach would seem to lend itself very well to the sort of inclusive and cross-domain representations that Wiggins and others have referred to and there are ongoing discussions to try and

[35] http://www-128.ibm.com/developerworks/java/library/j-camusic/#N102C3
[36] http://www.visualizationsoftware.com/
[37] As reported in Wired magazine, 2002
http://www.wired.com/news/culture/0,1284,52426,00.html

define exactly what might be required by the definitive musical critical edition model. Wiering, Crawford and Lewis have proposed a 'multidimensional model' that they argue solves the problem of current editorial methods,

> …by giving full access to as many source representations as one needs, by defining an edition as an adaptive layer on top of the sources, and by offering flexible generation and presentation of textcritical information.[38]

The preparation and presentation of materials for this type of multimedia production might involve quite a number of tools and techniques: from optical musical recognition tools such as Sharpeye,[39] for the capture and encoding of music notation into MIDI and MusicXML formats; to high end applied analysis tools based on the MPEG-7 or Humdrum toolsets. Bearing in mind the potential for the inclusion of related textual and image material, there may well also be scope for the kind of text analysis tools familiar to literary and linguistics scholars[40] or the methods used by the DIAMM project[41] when dealing with representations of original scores.

In the area of computer science, perhaps one of the most interesting developments of recent years, and one that is bound to have an impact on music and musicology, is the recent focus on the uses of Grid computing to arts and humanities researchers. The stated aim of this type of computing methodology is to enable greater storage and processing capacity to be brought to bear on data that is currently problematic due to its size and complexity, and this perfectly describes the scenario that Casey outlines when introducing the MP7AR tool described above.

> Finding similar passages in a large collection requires an enormous amount of computation. For collections greater than a few works, this computation becomes intractable because of the combinatorial bottleneck that grows rapidly as the database becomes large.[42]

In a poster for the Euroweb 2002 conference, Matthew Dovey outlined the benefits of setting up Grid retrieval capabilities across virtual organisations[43] building on projects such as OMRAS[44] which set out to provide modular frameworks for MIR activities. The availability of distributed tools and resources is described as a Web Services model and this approach is now beginning to be discussed more widely in arts and humanities research. The distributed, highly resourced and secure model that Grid computing supplies potentially provides a way forward for various problems faced by musicology researchers, not only in areas of data storage and processing, but also in the aggregation of small datasets of distributed material (of which there are many), and the problems associated with accessibility and rights management. If this model gains widespread acceptance and new areas of analysis and research become available to large numbers of researchers, the use of ICT methods in musicology may begin to routinely represent genuinely new research rather than simply being a transposition of old research questions in the guise of technological presentation.

Neil Grindley,
Senior Project Officer
Methods Network
Version control – 3 July 2007

---

[38] Wiering et al (2006)
[39] http://www.visiv.co.uk/
[40] http://users.ox.ac.uk/~ctitext2/enquiry/tat01a.html
[41] http://www.diamm.ac.uk/
[42] Casey (2006)
[43] Dovey (2002)
[44] Online Music Recognition and Searching http://www.omras.org/Full_desc.html

References

Babbitt, M., 'The Use of Computers in Musicological Research', *Perspectives of New Music,* 3(2):74-83, 1965. Available at http://www.jstor.org/.

Bonardi, A., *What the Musicologist Needs*, International Symposium on Music Information Retrieval, 2000 http://ismir2000.ismir.net/papers/invites/bonardi_invite.pdf

Bresson, J., Agon, C., Assayag, G., *Open Music 5: A Cross-Platform Release of the Computer-Assisted Composition Environment*, IRCAM – Centre Pompidou article, 2005 http://mediatheque.ircam.fr/articles/textes/Bresson05b/#search=%22bresson%20openmusic%205%22

Casey, M., *Audio Tools for Music Discovery*, Methods Network Expert Seminar: Modern Methods for Musicology, Royal Holloway, University of London, 3/3/2006. http://www.methodsnetwork.ac.uk/redist/pdf/casey.pdf

Cook, N., *Keynote Talk: Towards the Compete Musicologist*, ISMIR 2005 conference proceedings http://ismir2005.ismir.net/proceedings/cook.pdf#search=%22OMRAS2%22

Crawford, T., Report on ISMIR 2002 Conference Panel II: OMRAS Approach to Music Information Retrieval in the Real World (2002) http://www.dlib.org/dlib/november02/11inbrief.html

Dovey, M., *GRID IR - GRID Information Retrieval*, Euroweb 2002 Conference
The Web and the GRID: from e-science to e-business, St Anne's College, Oxford, UK, 17-18/12/2002. http://www.w3c.rl.ac.uk/Euroweb/poster/112/gridir.html

Eakins, J., Graham, M., *Content-based Image Retrieval*, Technical Report JTAP039, JISC Technology Application Program, 2000

Eidenberger, H., *A New Perspective on Visual Information Retrieval*, Vienna University of Technology, 2004, http://www.ims.tuwien.ac.at/~hme/papers/ei2004-vir.pdf

Eisenberg, G., Lerch, A., Tanghe, K., FEAPI: A Low Level Feature Extraction Plugin API, Proceedings of the 8th International Conference on Digital Audio Effects, Madrid, Spain, 20-22//8/2005 http://feapi.sourceforge.net/

Fingerhut, M., Donin, N., Filling Gaps Between Current Musicological Practice and Computer Technology at IRCAM, Methods Network Expert Seminar: Modern Methods for Musicology, Royal Holloway, University of London, 3/3/2006. http://www.methodsnetwork.ac.uk/redist/pdf/fingerhut.pdf

Fujinaga, I., Forscher Weiss, S., 'Music', in Schreibman, S., Siemens, R., Unsworth, J., (eds), *A companion to Digital Humanities*, (pp. 97 – 107)

Jeffrey, S., *Life Beyond Google,* Paper given at DRHA 2006, Dartington College of Arts http://www.dartington.ac.uk/drha06/papers/abstract.asp?uid=47

Lindsay, A., *Understanding the Capabilities of ICT Tools for Searching, Annotation and Analysis of Audio-Visual Media*, Methods Network Expert Seminar: Modern Methods for Musicology, Royal Holloway, University of London, 03/03/2006 http://www.methodsnetwork.ac.uk/redist/pdf/lindsay.pdf

Roland, P., The Music Encoding Initiative(MEI) DTD and the OCVE, University of Virginia Library http://www.ocve.org.uk/redist/pdf/roland.pdf#search=%22ocve%20mei%22

Wiering, F., Crawford, T., Lewis, D., *Digital Critical Editions of Music: A Multidimensional Model*, Methods Network Expert Seminar: Modern Methods for Musicology, Royal Holloway, University of London, 3/3/2006. http://www.methodsnetwork.ac.uk/redist/pdf/wiering.pdf

Wiggins, G., *Computer Representation of Music in the Research Environment,* Methods Network Expert Seminar: Modern Methods for Musicology, Royal Holloway, University of London, 3/3/2006. http://www.methodsnetwork.ac.uk/redist/pdf/wiggins.pdf

Further Resources and Useful Links

CHARM
http://www.charm.rhul.ac.uk/content/purpose/purpose.html
AHRC funded project focusing on recording and performance in musicology

Digital Music Research Network List
http://www.jiscmail.ac.uk/lists/DMRN-LIST.html
Mailing list

Grove Music Online
http://www.grovemusic.com/shared/views/article.html?section=music.40583.2
Computers and Music

HUMDRUM toolkit
http://dactyl.som.ohio-state.edu/Humdrum/guide01.html
Comprehensive information on tools suite

IRCAM
http://www.ircam.fr/61.html
Centre Pompdou – french language website of leading Music research centre

MELDEX
http://www.dlib.org/dlib/may97/meldex/05witten.html
The New Zealand Digital Library MELody index – query by humming

MuseData
http://www.musedata.org/
Data Encoding format

Music Notation
http://www.music-notation.info/
Very rich resource focused on music notation and data transfer issues

OMRAS 2
http://www.omras.org/Full_desc.html
Online Music Recognition and Searching

Variations 2
http://variations2.indiana.edu/

Large digital sound library initiative