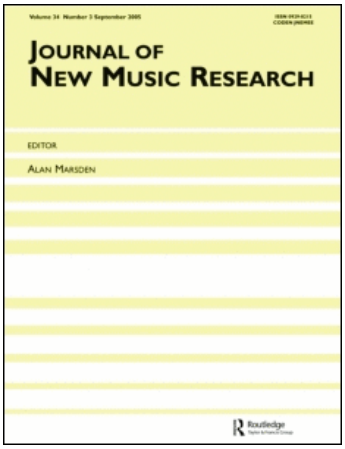

Journal of New Music Research
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713817838

Towards a Distributed Research Environment for Music Informatics and Computational Musicology

Simon Dixon[a]; Mark Sandler[a]; Mark d'Inverno[b]; Christophe Rhodes[b]
[a] Queen Mary University of London, UK [b] Goldsmiths University of London, UK

Online publication date: 19 February 2011

To cite this Article Dixon, Simon , Sandler, Mark , d'Inverno, Mark and Rhodes, Christophe(2010) 'Towards a Distributed Research Environment for Music Informatics and Computational Musicology', Journal of New Music Research, 39: 4, 291 — 294
To link to this Article: DOI: 10.1080/09298215.2010.545241
URL: http://dx.doi.org/10.1080/09298215.2010.545241

Journal of New Music Research
2010, Vol. 39, No. 4, pp. 291–294

Routledge
Taylor & Francis Group

Towards a Distributed Research Environment for Music Informatics and Computational Musicology

Simon Dixon[1], Mark Sandler[1], Mark d'Inverno[2], and Christophe Rhodes[2]

[1]Queen Mary University of London, UK; [2]Goldsmiths University of London, UK

Welcome to this special issue of the *Journal of New Music Research*, on the OMRAS2 project: Online Music Recognition And Searching—A Distributed Research Environment for Music Informatics and Computational Musicology. This article introduces the six papers in this issue, setting them in the broader context of music informatics and computational musicology research, and gives a brief overview of other work completed in the OMRAS2 project. This is followed by reviews of the achievements of the OMRAS2 project from the principle investigators at the two institutions where the project took place: Queen Mary University of London and Goldsmiths University of London.

Musicology and informatics, once unrelated fields, now share considerable common ground. On the one hand, computer software for the analysis of musical content in score or audio formats is winning an important place in empirical (or systematic) musicology (Cook, 2004; Crawford and Gibson, 2009), for example where it facilitates the analysis of large music collections (Widmer, 2001). On the other hand, the digitisation of the music production, distribution and consumption chain, fuelled by rapid advances in affordable data processing, storage and transmission, has led to the availability of vast quantities of musical content and metadata via the Internet. These developments have overturned traditional business models for the music industry and at the same time spawned the field of music informatics, also known as Music Information Retrieval, which deals with the automation of processes associated with accessing, analysing and manipulating musical content and knowledge.

Most research in music informatics has focussed on the development of software components, which only rarely are integrated into complete systems, thereby limiting their usefulness for researchers from other fields. The OMRAS2 project proposed to address this issue by developing an integrated web-based platform to provide tools and infrastructure supporting computational musicology.

The first two papers in the issue address the use of the Semantic Web in computational musicology and music informatics research. Fazekas, Raimond, Jacobson, and Sandler (2010) explain the Semantic Web vision of creating compatible software via data compatibility, and they describe the ontologies, data and tools that have been developed in order to lay the groundwork for realising this vision within the field of computational musicology. Cannam, Sandler, Jewell, Rhodes, and d'Inverno (2010) take a critical look at current Semantic Web technologies, assess the extent to which they are able to facilitate distributed music informatics research, and discuss open problems which need to be solved in order to achieve the goal of seamless interoperable software. The third paper in the collection describes a musicological application combining signal processing and Semantic Web tools in order to estimate the temperament of a harpsichord from a recording of any standard solo work (Tidhar, Fazekas, Mauch, & Dixon, 2010). The application is a publicly available web service which uses an extensible temperament ontology to model musical temperament and facilitate the description and inference of the characteristics of specific temperaments.

Rhodes, Crawford, Casey, and d'Inverno (2010) present the audioDB software for content-based similarity search, and examine several use cases covering a range of different music collection sizes, showing that the

Correspondence: Simon Dixon, School of Electronic Engineering and Computer Science, Queen Mary University of London, Mile End Rd, London, E1 4NS, UK. E-mail: simon.dixon@eecs.qmul.ac.uk

DOI: 10.1080/09298215.2010.545241

optimal balance of precision and recall is dependent on the task and data set at hand. In the fifth paper, Anglade, Benetos, Mauch, and Dixon (2010) present research on automatic characterisation of music by its harmonic progressions, applied to the task of genre classification. The addition of high-level harmony features to a standard timbre-based genre classification framework is shown to improve classification accuracy. In the final paper Flossmann, Goebl, Grachten, Niedermayer, and Widmer (2010) describe a large-scale computational musicology project analysing expressive performance in recordings of the complete works of Chopin for solo piano, performed by the Russian pianist Nikita Magaloff. This provides a perspective of computational musicology research independent of and complementary to the OMRAS2 project.

This special issue is by no means complete in its coverage of the OMRAS2 project. Much OMRAS2 research is published elsewhere, including work on user interfaces (Magas, Casey, & Rhodes, 2008), standardisation (Jang, Kudumakis, Sandler, & Kang, 2011), sinusoidal modelling (Wen and Sandler, 2010a, 2010b), chord transcription (Mauch and Dixon, 2010), structural segmentation (Levy and Sandler, 2008), key detection (Noland and Sandler, 2009), audio effects (Wilmering and Sandler, 2010), playlist generation (Fields, Jacobson, Rhodes, & Casey, 2008), score following (Macrae and Dixon, 2010; Arzt, Widmer, & Dixon, 2008), style characterisation (Jewell, Rhodes, & d'Inverno, 2010) and ethnomusicology (Proutskova and Casey, 2009).

—Simon Dixon, Guest Editor

Queen Mary

This special issue has been put together to present some of the outcomes of the EPSRC (Engineering and Physical Science Research Council) funded project known as OMRAS2. The OMRAS2 project set out to follow in the footsteps of the trail-blazing OMRAS project (`www.omras.org`), which brought together partners from the USA and the UK and was funded by JISC (the Joint Infrastructure Sub-Committee of the Higher Education Funding Council of England—HEFCE) and the National Science Foundation (NSF). That project has been described in several papers, the most comprehensive of which is by Pickens et al. (2003).

Those of us involved in OMRAS and OMRAS2 have almost forgotten what the acronym stands for! It is Online Music Recognition and Searching, though one OMRAS2 team member notably re-purposed the acronym as Ontology-driven Music Retrieval and Annotation Sharing service—which in many senses it is.

The achievements of the project are summarised on its web-site, `www.omras2.org`, which includes resources to learn more about ontologies, especially the Music Ontology, which the project is proud to have founded, and the Semantic Web.

While one of the project's initial aims was to investigate the use of Web Services and a Service Oriented Architecture for Music Informatics, this aspiration was quickly superceded by the more general and significantly more powerful Resource Oriented Architecture, based on Resource Description Framework (RDF), Ontologies, the Semantic Web and Open Linked Data. Many of our early project meetings were devoted to discussing which communications architecture we should build our infrastructure around, but all this was swept away once we understood the deeply powerful paradigm that is the Semantic Web. I for one have come to understand Ontologies and RDF as a powerful new technology, not merely a representation scheme. This special issue will, I hope, help others to 'get' the SemanticWeb and start to adopt it in their own projects, products and services. One really telling point is that the original prototype of SAWA (Sonic Annotator Web Application) was built in a week by a single (albeit very gifted) researcher/programmer.

While this special issue highlights some of the project's achievements, there simply isn't room to cover them all, and the interested reader is invited to look at the web site (see above). Here one can find (at the time of writing) over 100 papers and presentations that have resulted. There is so far one completed PhD thesis, with another four in the final stages of writing. Several other PhD students have been associated with the project and have both contributed to it, and benefited from the research and the resources that have resulted. More OMRAS2-related PhD theses will result. Hopefully there are even several OMRAS2-influenced PhDs from universities other than Goldsmiths and Queen Mary: that was one of our original aims.

There were nine workshops. The first was held at Queen Mary in December 2008, immediately after the Centre for Digital Music's annual Digital Music Research Network Workshop. Others—all listed on `www.omras2.org/Workshops`—have been held at the ISMIR conference series, the ACM Recommender Systems Conference, and at 2 Music Hack Days. Well in excess of 1000 researchers have been at one or more of these workshops.

What is next? Several other projects have started up. In one, some of the team is working with the British Library Sound Archive and some UK secondary schools to promote 'Musicology for the Masses' (`gow.epsrc.ac.uk/ViewGrant.aspx?GrantRef=EP/I001832/1`). In another, we are working with MusicBrainz to fully integrate their metadata services with the Semantic Web, funded by JISC once more (`linkedbrainz.c4dmpresents.org`). In another, Queen Mary and Goldsmiths are partnered by colleagues in Southampton and Oxford Universities, McGill University in Canada,

and Waikato in New Zealand under the leadership of Professor Stephen Downie of the University of Illinois at Urbana-Champaign: this is funded by the Mellon Foundation as the NEMA Project (Networked Environment for Music Analysis—www.nema.lis.uiuc.edu), and at the time of writing there are two more industrially focussed proposals awaiting funding decisions. We expect more proposals and grants to be forthcoming.

—Mark Sandler
Queen Mary University of London

Goldsmiths

OMRAS2 was a joint project between Queen Mary and Goldsmiths with stakeholders including King's College London, Royal Holloway University of London, Yahoo, and AWAL (Artists Without a Label). The main focus of the project is around information retrieval in the field of music but encompasses work in databases and knowledge representation, statistics, design and traditional musicology as well as technological research and development.

Central to our vision is enabling interested parties to discover new information about musical content they have access to whether that is a personal collection, an academic archive or a commercial catalogue for purchase or sale. Part of the challenge here as we see it is in marrying, in a usable form, diverse sources of information: not just from the audio content itself but from metadata, social networking information and free and structured text related to music that increasingly predominates on the wider internet. The integration of these sources of information allows our technology to serve a variety of different use cases with varied multimedia in different forms.

Because we are deploying tools and techniques for use by non-technical users, the input of designers is critical to provide powerful and intuitive visual entry points into this extremely complex world of search and retrieval. Three user interfaces developed during OMRAS2 (Magas et al., 2008; Magas and Proutskova, 2009; Magas, Stewart, & Fields, 2009) illustrate the designer's perspective on making technological research of this kind accessible, by harnessing metaphor to provide the user with an intuition of what is going on under the hood, without exposing the underlying complexity that can only be appreciated after extensive experience or training.

In order to provide the information for our user interfaces it is necessary for us to decide on an interchange format and in recent years the semantic web has developed the potential to render information in many forms both accessible to other computer agents such as user interfaces and inference engines as well as to the end user. We have therefore adopted the semantic web paradigm for our data and knowledge representation and Cannam et al. (2010) describe both the successes and the areas where more work is needed in the representation of music informatic knowledge and the adaptation and construction of software tools developed in the OMRAS2 project in collaboration between Queen Mary and Goldsmiths.

One of the most exciting areas of research in this field is the generation of new knowledge and insights into content. The audioDB system described by Rhodes et al. (2010) allows its users to search for relationships between content—looking at associations ranging from fragments of audio, through whole tracks to large commercial catalogues. This covers tasks such as grouping recordings by sonic similarity, identifying performances of the same work (or cover versions), through to deciding whether a database contains items which use fragments of copyright material from other sources.

As well as the work presented in this issue, we have applied our research to ethnomusicological cataloguing, live coding performance, exploration of images and film, and indexing and retrieval of motion capture data. A lot of the tools and research form this project will continue to be developed in the context of NEMA (Networked Environment for Musical Analysis), a collaborative enterprise between a number of music informatics research centres (see www.nema.lis.uiuc.edu).

—Mark d'Inverno and Christophe Rhodes
Goldsmiths University of London

Acknowledgement

This work was supported by the EPSRC project OMRAS2 (EP/E017614/1).

References

Anglade, A., Benetos, E., Mauch, M., & Dixon, S. (2010). Improving music genre classification using automatically induced harmony rules. *Journal of New Music Research*, *39*(4), 349–361.

Arzt, A., Widmer, G., & Dixon, S. (2008). Automatic page turning for musicians via real-time machine listening. In M. Ghallab, C. Spyropoulos, N. Fakotakis, & N. Avouris (Eds.), *Proceedings of the 18th European Conference on Artificial Intelligence (ECAI 2008)* (pp. 241–245). Amsterdam: IOS Press.

Cannam, C., Sandler, M., Jewell, M.O., Rhodes, C., & d'Inverno, M. (2010). Linked Data and you: Bringing music research software into the Semantic Web. *Journal of New Music Research*, *39*(4), 313–325.

Cook, N. (2004). Computational and comparative musicology. In E. Clarke & N. Cook (Eds.), *Empirical Musicology: Aims, Methods, Prospects* (pp. 103–126). New York: Oxford University Press.

Crawford, T., & Gibson, G. (Eds.) (2009). *Modern Methods for Musicology: Prospects, Proposals, and Realities*. Digital research in the arts and humanities. Farnham: Ashgate.

Fazekas, G., Raimond, Y., Jacobson, K., and Sandler, M. (2010). An overview of Semantic Web activities in the OMRAS2 project. *Journal of New Music Research*, *39*(4), 295–311.

Fields, B., Jacobson, K., Rhodes, C., & Casey, M. (2008). Social playlists and bottleneck measurements: Exploiting musician social graphs using content-based dissimilarity and pairwise maximum flow values. In *International Conference on Music Information Retrieval*, pp. 559–564.

Flossmann, S., Goebl, W., Grachten, M., Niedermayer, B., & Widmer, G. (2010). The Magaloff project: An interim report. *Journal of New Music Research*, *39*(4), 363–377.

Jang, I., Kudumakis, P., Sandler, M., & Kang, K. (2011). The MPEG interactive music application format standard. *IEEE Signal Processing Magazine*, *28*(1), to appear.

Jewell, M. Rhodes, C., & d'Inverno, M. (2010). Querying improvised music: Do you sound like yourself? In *11th International Society for Music Information Retrieval Conference*, Utrecht, The Netherlands, pp. 483–488.

Levy, M., & Sandler, M.B. (2008). Stuctural segmentation of musical audio by constrained clustering. *IEEE Transactions on Audio, Speech and Language Processing*, *16*(2), 318–326.

Macrae, R., & Dixon, S. (2010). Accurate real-time windowed time warping. In *11th International Society for Music Information Retrieval Conference*, Utrecht, The Netherlands, pp. 423–428.

Magas, M., Casey, M., & Rhodes, C. (2008). mHashup: Fast visual music discovery via locality sensitive hashing. In *International Conference on Computer Graphics and Interactive Techniques (ACM SIGGRAPH) New Tech Demos*, Los Angeles, CA, article number 26.

Magas, M., & Proutskova, P. (2009). A location-tracking interface for ethnomusicological collections. In *European Conference on Digital Libraries, Workshop on Exploring Musical Information Spaces*, Corfu, Greece, pp. 12–17.

Magas, M., Stewart, R., & Fields, B. (2009). decibel 151: Collaborative spatial audio interactive environment. In *International Conference on Computer Graphics and Interactive Techniques (ACM SIGGRAPH) Art Gallery*, New Orleans, LA, article number 21.

Mauch, M., & Dixon, S. (2010). Simultaneous estimation of chords and musical context from audio. *IEEE Transactions on Audio, Speech and Language Processing*, *18*(6), 1280–1289.

Noland, K., & Sandler, M. (2009). Influences of signal processing, tone profiles, and chord progressions on a model for estimating the musical key from audio. *Computer Music Journal*, *33*(1), 42–56.

Pickens, J., Bello, J., Monti, G., Sandler, M., Crawford, T., Dovey, M., & Byrd, D. (2003). Polyphonic score retrieval using polyphonic audio queries: A harmonic modelling approach. *Journal of New Music Research*, *32*(2), 223–236.

Proutskova, P., & Casey, M. (2009). You call that singing? Ensemble classification for multi-cultural collections of music recordings. In *10th International Society for Music Information Retrieval Conference*, Kobe, Japan, pp. 759–764.

Rhodes, C., Crawford, T., Casey, M., & d'Inverno, M. (2010). Investigating music collections at different scales with audioDB. *Journal of New Music Research*, *39*(4), 337–348.

Tidhar, D., Fazekas, G., Mauch, M., & Dixon, S. (2010). TempEst: Harpsichord temperament estimation in a Semantic Web environment. *Journal of New Music Research*, *39*(4), 327–336.

Wen, X., & Sandler, M. (2010a). On the characterization of slowly varying sinusoids. *EURASIP Journal on Audio, Speech, and Music Processing, 2010*, article ID 941732.

Wen, X., & Sandler, M. (2010b). Source-filter modeling in sinusoid domain. *Journal of the Audio Engineering Society*, *58*(10), 795–808.

Widmer, G. (2001). Using AI and machine learning to study expressive music performance: Project survey and first report. *AI Communications*, *14*(3), 149–162.

Wilmering, T., & Sandler, M. (2010). RDFx: Audio effects utilising musical metadata. In *IEEE International Conference on Semantic Computing*, Pittsburgh, PA, pp. 452–453.