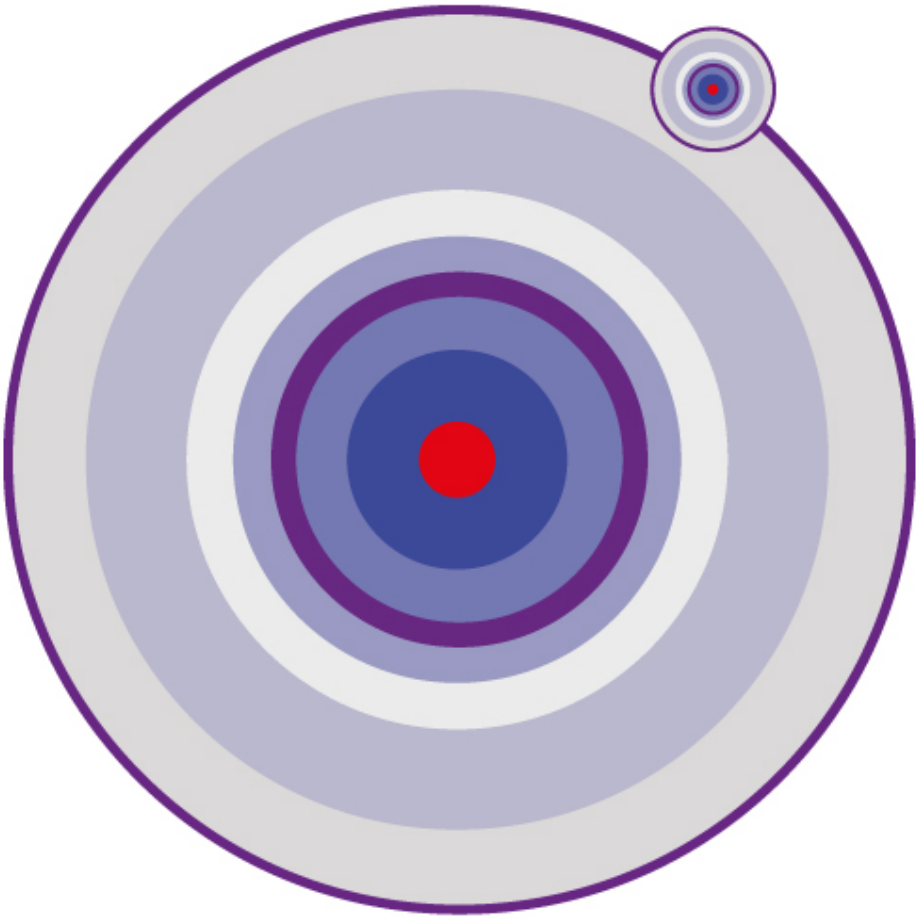

MIRES

PUBLIC REPORT

SEVENTH FRAMEWORK
PROGRAMME

TABLE OF CONTENTS

1 BACKGROUND

The Roadmap for Music Information ReSearch (MIReS) is a Coordination and Support Action funded by the 7th Framework Programme of the European Commission and coordinated by the Music Technology Group at Universitat Pompeu Fabra in Barcelona in collaboration with the London design research company Stromatolite LTD. The project was launched on the 1st of October 2011 and runs for 18 months. Partners include the Austrian Research Institute for Artificial Intelligence (OFAI) in Vienna, Institut de Recherche et Coordination Acoustique/Musique (IRCAM) at the Centre Pompidou in Paris, Institute for Systems and Computer Engineering of Porto (INESCP), Centre for Digital Music at Queen Mary University of London (QMUL) and Barcelona Music and Audio Technologies (BMAT).

2 INTRODUCTION

The MIReS project aims to create a research roadmap of the MIR field, by expanding its context and addressing challenges such as multimodal information, multiculturalism and multidisciplinarity. MIR has the potential for a major impact on the future economy, the arts and education, not merely through applications of technical components, but also by evolving to address questions of fundamental human understanding, with a view to building a digital economy founded on "uncopiable intangibles": personalisation, interpretation, embodiment, findability and community. Within this wider context we propose to refer to the field of MIR as Music Information ReSearch (MIReS) and thus widen its scope, ensuring its focus is centered on quality of experience with greater relevance to human networks and communities.

3 PROJECT OBJECTIVES

The field of Music Information Retrieval (MIR) has centred primarily on the analysis of sound for the purpose of more efficient search and faster access to digital collections of recorded music. The advent of web-mediated social networks has led to the evolution of music networks and numerous music-related communities. This has created a dynamic global market for digital music and collateral products and services with significant challenges and huge opportunities for exploitation. For Europe to leverage its position as a world leader in music creativity, production and mobile distribution, a programme of digital music search technology is needed in order to ensure coherent targeted support for innovation and underpin competitive strategies for maintaining European excellence. The roadmap for the future of MIR includes a framework for cooperation and co-creation across academic communities and a virtual centre of excellence for music-related and MIR-relevant studies.

The Roadmap for MIReS aims to provide a meta-analysis of the MIR discipline, address emergent contexts and major challenges, formulate research evaluation standards for the discipline, contribute to the establishment of music production and digital library management standards, engage a variety of stakeholders from different disciplines of academia and industry and deliver innovative platforms for co-creative workshops focusing on horizon-scanning and technology foresight. The roadmap also aims to include a framework for the establishment of a MIR excellence network, involving the field drivers and stakeholders (music researchers, industry representatives, strategists and policy makers).

3.1 Objectives

- Leverage European leadership in music creativity, production and mobile distribution
- Build a digital economy founded on ‘uncopiable intangibles’: personalisation, interpretation, embodiment, findability, community
- Ensure coherent targeted support for innovation to underpin a competitive strategy for the European music industry and the future of music search
- Widen and deepen MIR - seen as Music Information ReSearch - based on music as a cultural construct with greater relevance to human networks and communities
- Provide support for multimodality, multiculturalism, multidisciplinarity, cognitive science and AI
- Enhance co-creativity in music search and discovery, and foster rich user Quality of Experience (QoE) enabling a multiplicity of interpretations and the emergence of new user behaviours

3.2 Methods

- Provide a meta-analysis of the discipline
- Examine applications of MIR
- Address emergent contexts
- Address major challenges
- Formulate research evaluation standards

1 Contribute to industry standards

- Deliver innovative platforms
- Involve a variety of stakeholders
- Engage companies and researchers from outside the EU
- Engage the music industry in debates and events
- Ensure dissemination in industry and establish long-term stakeholder relationships
- Build a framework for a virtual network of MIR excellence

1.1 Structure

The overall strategy of the work within the scopes of the project has been divided into distinct 6 Work Packages, each addressing an essential component area of the Coordination Action.

WP 1 Project Management

Work Package 1 outlines a plan for effective management of the Coordination Action and the related reporting strategy.

WP 2 Meta-Analysis of the MIR Discipline

Work Package 2 focuses on the Meta-Analysis of the MIR Discipline, and is divided into distinct thematic areas, each one handled by a consortium partner. It has direct influence over the content used in Dissemination (WP4) and in laying the foundations of the writing of the Roadmap for Music Information ReSearch (WP3).

WP 3 Roadmap Document for Music Information ReSearch

Work Package 3 focuses on creating the Roadmap for Music Information ReSearch covering outstanding research issues, horizon scanning & technology foresight and market futures for related services. It is built upon the report generated by WP2 in Month 9 and on the knowledge gathered and generated during the events organised by Community Co-creativity and New Knowledge Generation: Hubs and Spokes (WP5).

WP 4 Dissemination: Wiki, Publications, Conferences and Workshops

Work Package 4 is dedicated to the wide dissemination of knowledge about MIReS as well as collecting information relevant for the roadmap, gathered both from the meta-analysis of the discipline (WP2) and from the Roadmap findings (WP3) and involves wiki, publications, conferences, workshops, international researchers and external stakeholders.

WP 5 Community Co-creativity and New Knowledge Generation (Hubs and Spokes)

Work Package 5 is devoted to Community Co-creativity and New Knowledge Generation: Hubs and Spokes, and involves members from academic disciplines, international researchers, different age groups, multicultural communities, events with EU Associate countries, and a variety of industry stakeholders.

WP 6 Framework for an MIReS Network of Excellence: Research-to-Industry

Work Package 6 provides the framework for a network of MIReS excellence from research-to-industry in view of ensuring efficient communication and transfer of innovation to industry in the long-term and continuing to inform policies beyond the lifetime of the proposal.

2 SUMMARY OF COMPLETED WORK

During the first 12 months of the 18-month project the majority of hubs and spokes were organised to generate new knowledge for the roadmap. The meta-analysis of the state of the art was completed and roadmapping challenges mapped.

Additional value was achieved by reaching a number of objectives which were planned for later in the programme. These included all of the methods listed under **Objectives**, thus creating a strong framework for dissemination (see D4.2), knowledge building for the roadmap (see D3.2) and a research-to-industry network (later reported in detail in deliverable D6.1).

The last six months of the project focused principally on the writing of the Roadmap, with feedback from the MIReS academic community and further input from stakeholders. Work has been devoted to securing platforms and events to run beyond the duration of the project thus ensuring long-term impact of knowledge generated by the Roadmap and related activities.

2.1 Meta-analysis of the state of the art (WP2)

During the first project meeting and roadmapping workshop in Barcelona in October 2011 it was concluded that the analysis of the state of the art could not be structurally isolated from the roadmap challenges as the two were interdependent. A document structure was therefore set to include both meta-analysis and roadmapping content. Topics covered therefore had to reflect both the main research areas of the state of the art, as well as areas which are not currently investigated but are considered important for future MIR research. For the latter, state of the art had to refer to evidence of research from other relevant current practices in the arts and sciences.

During the second roadmapping workshop in Vienna (2-3 February 2012) the overall roadmap structure was further improved and subject areas defined for the state of the art. Roadmap challenges were introduced into completed state of the art sections of the roadmap in the third roadmapping workshop in Barcelona (12-13 April 2012) and were further defined with evidence from organised events in the fourth workshop in London in June 2012. A summary of the state of the art achievements is available in the deliverable D2.2 and on the Meta-Analysis pages of the MIReS research wiki.

2.2 Roadmap challenges (WP3)

During the roadmapping process the roadmap challenges were being continuously updated and comments and contributions sought from the wider MIReS community.

The Roadmap kick-off meeting in Barcelona identified expertise and specific resources of the individual consortium partners which led to the assignment of specific roles and topic categories for the roadmapping process. At the roadmapping workshop in Vienna (2-3 February 2012) a concrete plan was identified for producing the roadmap which included a finalization of the overall roadmap structure and a concrete plan for the actual compilation and writing of the roadmap.

The aim of the following meeting in Barcelona (12-13 April 2012) was to review all the material already collected for the "state -of-the-art" part of the roadmap and to discuss requirements to further improve this part of the document. In addition specific roles and tasks for the "challenges" part of the roadmap were assigned. At the London meeting (18 June 2012) all text and information concerning the "challenges" part of the roadmap was reviewed and requirements to further improve the "challenges" part of the document were identified.

Prior to the ISMIR 2012 conference (8-12 October 2012) a preliminary version of the roadmap was published on the MIReS Wiki pages to allow feedback from the MIR community. This was announced in due time before the ISMIR conference via the MIR mailing list with the aim to stimulate discussion at the main conference for MIR research. Interested researchers were able to submit comments directly to the Wiki page. A questionnaire was also published about selected topics and challenges from the roadmap via the MIR mailing list. Participants were able to rank each challenge on a scale from 1 (not relevant) to 5 (very relevant) and to give direct feedback via submission of comments. This allowed to get further feedback prior to the ISMIR conference and enabled to identify topics and aspects which had previously been overlooked.

During the ISMIR 2012 conference feedback was gathered via a series of sessions sponsored and organised by the MIReS consortium, including:

- the MIRrors session (MIRrors: Looking back to the past of ISMIR to face the future of MIR) where authors were asked to present reflections on the evolution of the MIR field or a particular MIR topic and provide insights on short-term future challenges.
- the Special Session on Evaluation Initiatives in MIR, including MIREX, the MillionSong data set and MusiClef.
- the MIReS late breaking session on "MIReS Roadmap: Challenges for Discussion" presented a selection of challenges from the roadmap. Members from the consortium were able to discuss these with all interested researchers attending. A short summary of this session has been documented at the ISMIR 2012 website.

Following ISMIR activities, the main goal of the MIReS workshop in Paris (8-9 November 2012) was the analysis, discussion and integration of all the feedback gathered through the dissemination of the intermediate version of the roadmap to the general public. Finalising of the content was subject of an intensive three-day workshop in Barcelona (22-24 January 2013). The MIReS Roadmap has since been formatted and made available for publication and dissemination via the mires.cc channel and the music_ir list. A shortened and condensed version of the final roadmap will be submitted to the journal of "ACM Computing Surveys". Publication in this journal will enable the circulation of the roadmapping results to a general computer science audience.

2.3 Dissemination of roadmapping efforts (WP4)

The project website and documentation hub were established in Month 1 of the project. The public-facing website (http://www.mires.cc) has been designed as a media hub collating videos of research work by the partners and videos of the roadmapping process. The documentation hub, linked to the main site, was made available as a MIReS research wiki containing all of the information about the roadmapping process.

The project has been disseminated at a variety of conferences and workshops addressing the MIR community, external related academic conferences, industry stakeholder conferences and events, EU concertation meetings and workshops. A full list of conferences and events is available on the MIReS wiki proceedings pages.

2.3.1 EU meetings and workshops

MIReS has contributed to the Networked Media 8th FP7 Concertation Meeting in Brussels on the 13-14 December 2011; the Workshop on the FP7 ICT Work Programme 2013 for the Networked Media sector in Brussels on the 27th of January 2012; the Search Computing workshop on the 25-26 September 2012; and the NEM Summit 2012 in Istanbul, Turkey.

2.3.2 MIR community conferences and events

Dissemination has been achieved at the yearly conference of the International Society for Music Information Retrieval (ISMIR 2012), and academic conferences focusing on music information and signal processing (FRSM 2012, ICASSP 2012, AdMIRe 2012, INFLA/FMA 2012, SMC 2012, DAFx 2012, ICMC, EUSIPCO).

2.3.3 Wider MIReS community academic conferences

In view of widening the scope of Music Information ReSearch dissemination has been achieved at conferences addressing MIReS challenges: TEI 2012 (multimodal

interaction), WWW 2012 (World Wide Web standards) NIME 2012 (musical interfaces), Manifeste Festival (multidisciplinarity), CMMR 2012 (music and emotion), CloudBusting 2012 (cloud computing), ICMPC 2012 (music cognition), ACM Multimedia 2012 (multimodality), CloudBusting 2012 (cloud computing), and MWC 2013 (mobile technologies).

2.3.4 Conferences involving industry and research

MIReS has been disseminated at MIDEM 2012 and 2013, Music Tech Fest 2012, Music Matters 2012, Music Hack Day Paris 2012, SONAR 2012, and the First Classical Music Hack Day in Vienna on 1-3 February 2013.

2.4 Knowledge-generating hubs and spokes

During the reporting period the consortium organised a series of specially conceived events which have successfully contributed to knowledge generation and broadening of scope for the area of Music Information ReSearch. Several of the events have attracted worldwide online audiences and an ongoing base of followers.

All of the MIReS-organised events have been fully documented and have generated valuable feedback for the roadmap. A more detailed summary of the events can be found in the deliverable D5.1. Videos from the specially organised events are available on the MIReS website, and on the individual video links in the wiki proceedings pages.

*2.4.1 Panel discussion: **Music Information Research of relevance to Hindustani and Carnatic music** at the 1st CompMusic Workshop at the KIIT-Gurgaon, India (20 January 2012)*

The aim of the panel was to talk about melodic and rhythmic characteristics of both Hindustani and Carnatic music, focusing on the differences between the two cultures in terms of intonation (tuning), melodic and motivic organization, and rhythmic patterns. Some ideas for the panel discussion included:

- differences between tabla and mridangam in terms of rhythmic playing
- improvisation styles based on melodic signatures and phrases in both Hindustani and Carnatic musics
- Gamakas in Carnatic music versus meends/gamakaas in Hindustani music
- characterization of Gharanas in Hindustani music in terms of intonation, melody, and rhythm
- types of singing Khayal vs. Drupap

2.4.2 Workshop: ***4th International Workshop on Advances in Music Information Research****, AdMIRe 2012, Lyon, France (17 April 2012)*

The workshop brought together world-class researchers in Music Information Research (MIR) to discuss topics highly relevant for the future of MIR. In particular, the following short-term objectives of MIReS have been addressed in AdMIRe:

• Formulate research evaluation standards (paper presentations by Urbano and by Bertin-Mahieux)

• Assess emerging contexts, such as web mining (paper presentation by Hauger)

• Engage researchers from outside the EU (paper presentations by Bertin-Mahieux and by Hankinson, participation of Fujinaga)

• Major challenges of MIR (multiculturalism in the keynote by Serra, semantic gap in the paper presentation by Sordo, multimodal information in the paper presentation by Hankinson, personalized and context-aware music retrieval and recommendation in the keynote by Ricci, scalability in the paper presentation by Bertin-Mahieux)

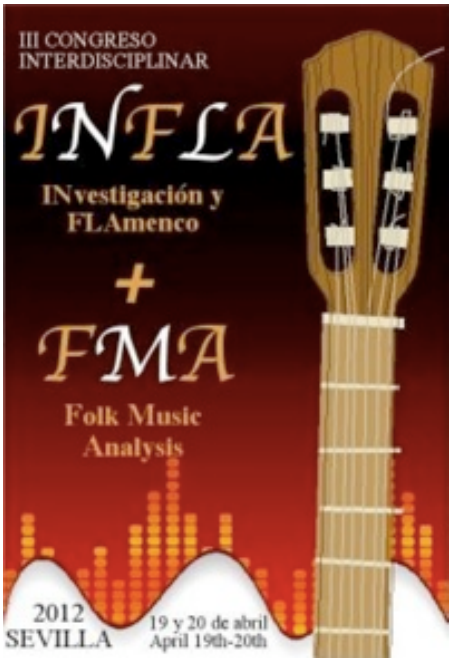

2.4.3 Panel discussion: ***Technological challenges for the computational modeling of the world's musical heritage*** *at the INFLA/FMA in Seville, Spain (20 April 2012)*

The following topics were discussed:

How to foster collaboration between researchers from different disciplines (musicologists, musicians, engineers...)?

- need to build truly interdisciplinary teams
- it is important to turn computational modes into usable tools that ethnomusicologists or musicologists can use
- need to define new methodologies adapted to the use of technology for ethnomusicological studies

Is it right to focus on single culture studies rather than approach musical phenomena cross-culturally?

- need to study single cultures but at the same time we need to keep links and collaborations between cultures to perform cross-cultural analyses.

Which should be our next steps in establishing ethno-Music Information Retrieval (ethno-MIR)?

- create an ethno-MIR Special Interest Group within the ISMIR society
- organize an annual meeting during ISMIR
- establish FMA as an annual forum for ethno-MIR

- write grant proposals for collaborations in ethno-MIR; in particular a proposal for a European COST action that would support conferences and workshops in ethno-MIR

*2.4.4 Talks, demos and performances at the launch of the **Music Tech Fest** with 52 contributors in London, UK (17-19 May 2012)*

The first **Music Tech Fest** gathered the entire music technology ecosystem under one roof – from the big brands of the music and media industries, music tech startups and apps creators, to developers, researchers, artists, performers, creatives and hackers. The event was co-funded by MIReS in conjunction with Ravensbourne's ERDF-funded Digital Media Innovation Consultancy and Business Incubator. Presentations, demos, installations, and performances from 52 contributors covered all aspects of music technology research and innovation within industry and academia.

Topics included: rights clearance, access to music collections for artists and researchers, music data visualisation, metadata creation and management, installation art driven by sound and music data, the impact of research on the music industry and innovative SMEs, new musical instruments driven by data, and performing with data. Written feedback and challenges have been gathered into a 45 page document with feedback and industry challenges for consideration by the MIReS consortium.

The festival attracted 1,036 attendees, an audience of over 700 live streaming viewers from 40 countries, and over 10,000 visitors to the site in May 2012 only. 830 minutes of video footage was captured on each of the three cameras and the 20 minute talks, demos and performances have been edited for the new Music Tech Fest: Music Tech Talks online media channel.

The Music Tech Fest video footage reached 656 different people over a period of two days, many of whom came back to view it again on several occasions, totaling 1435

views. Through live video streaming the Music Tech Fest reached 41 different countries, including Hong Kong, Greece, Russia, the US and Estonia. Since launching in London in May 2012, The Music Tech Fest has gained over 982 followers on Twitter, 436 of whom are active followers. The YouTube video release programme has enabled the Music Tech Fest to continue building a global following and to reach new territories like Brazil, that had not originally tuned in to the live broadcast during the festival.

The festival has established the foundation for regular Music Tech Talks. A follow-up event of Music Tech Talks was organised on July 26th 2012 as part of the Cultural Olympiad with contributions from music and creative industries (see 4.3.10).

Speakers included: Frederic Rousseau (IRCAM, research - MIReS partner), Christian Blom (artist), Matthew Davies (INESC - research, MIReS partner), Oscar Paytuvi (BMAT - SME, MIReS partner), Avi Ashkenazi (artist), Adam Place (Alphasphere - artist), Bruno Zamborlin (Mogees - artist and IRCAM/Goldsmiths research), Carles Lòpez (Reactable - artist), Matt Balck (Ninjatune - music industry, and Coldcut - artist), Nicole Yershon (Ogilvy Digital Labs - media industry), Estefania Caño (Fraunhofer / Songs2See - research), Saoirse Finn (QMUL - research, MIReS partner), Patrick Bergel (Animal Systems - SME, research), Ben Lawrence (Mixcloud - SME), Evan Stein (Decibel - SME), Kim de Ruiter (Noise Inc - SME), Matthew Sherett (Last.fm - SME), Tim Hadley (rara.com / Omnifone - SME), Philippe Perreaux (Right Clearing, SME), Cliff Fluet (Lewis Silkin - legal, music industry), Ed Averdieck (Cue Songs - SME), Will Page (PRS - music industry), Michela Magas (Stromatolite - SME and research, MIReS partner), Daniel Lewington (MPme / Apsmart - SME), Michael Breidenbrücker (RjDj - SME), Peter Kirn (Create Digital Music - research), Martin Ware (Illustrious - SME and music industry), Tom Cheshire (Wired - media industry), Jason Titus (CTO Shazam - music industry), Dave Haynes (VT Soundcloud - SME and music industry), DJ Ham (Ninjatune - music industry), Martin Macmillan (Soniqplay - SME), Paul D (artist), Olivier de Simone (webdoc - SME), Johann Waldherr (Spectral Mind - SME), Stephen O'Reilly (Mobile Roadie - SME and music industry), Ariel Elkin (London Music Hackspace - research), Jake Williams (artist and Goldsmiths music research), Daniel Jones (artist and Goldsmiths music research), Cassiel (artist), Jason Singh (artist / Victoria and Albert museum resident sound artist and music researcher), Matthew Hawn, Head of Product, Last.fm, Andrew Shoben, sound installation artist, Greyworld.

2.4.5 Workshop: ***Synaesthesia*** *at the Music Tech Fest in London. UK (17 May 2012)*

Ideas on "seeing music" we're used as a trigger for an investigation into the relationship between music data and the visual arts. Peter Kirn, author of Create Digital Music (createdigitalmusc.net) and of Synaesthesia workshops at Parsons School of Art in New York was invited to conduct the workshop. Tutorials were focused on methods in Processing and digital media, and examples drawn from digital ways of visualising music. The outcome showed a strong correlation between the interpretation of music and visual triggers. A collaborative application translating visual data to music data was used to generate music by drawing.

The final results were presented on the main stage to festival audiences and have been uploaded on the new Music Tech Fest / Music Tech Talks channel on YouTube.

*2.4.6 Workshop: **3D Music Hack Camp**, Music Tech Fest in London. UK (18-19 May 2012)*

The **3D Music Hack Camp** was conceived with the aim of combining the Internet of Things with music information hacking. The concept attracted contributions from some of the best known brands in the media and music industry: EMI, the BBC, Warp, Decibel, BMAT, Cisco, Ninja Tune, Last.fm and Animal Systems. The camp was run by Ariel Elkin from the London Music Hack Space.

The event generated novel applications of music information to objects like light sabres, play environments like ping pong tables, as well as new musical instruments and tangible interfaces. The final results were presented on the main stage to festival audiences and have been filmed for upload on the new Music Tech Fest / Music Tech Talks channel on YouTube.

2.4.7 Workshop: ***MIR and Creation*** *at the Manifeste Festival in Paris, France (2 June 2012)*

The goal of the workshop was to better understand the relationship between "M.I.R. and Creation" and study how MIR has extended its scope and is now used for the creation process itself. The invited key-speakers gave their point-of-view on the present and future of MIR for creation, being at the composition, interaction, performance or research level, at the audio, symbolic and database level. The main focus was on the use of M.I.R. for Creation, to go beyond "musaicing", and answer the questions:

- What can we do with the Information a-la-ISMIR (audio-descriptors, source-separation, chords, beats, auto-tags) for Creation?
- What other Musical Information can be used for Creation (symbolic, sensors, …)?

The goal of this workshop was to invite key actors to give their view on the use of M.I.R. technologies for creation. The goal was to answer the question - what uses for creation can we do with the Information a-la-ISMIR (audio-descriptors, source-separation, chords, beats, auto-tags)? - what other Musical Information could be used for creation (symbolic, sensors …) ? The speakers were invited according to their skills and representativity of a given field: representing research on music technologies (Assayag, Schnell, Schwarz, Jorda), composition (Manoury) or industry (Jehan,

Pachet). The workshop was made up of 7 keynotes of 1 hour each followed by a round-table of 1 hour.

Gerard Assayag mentioned the early used of audio-content information for music creation such as used in the spectral approach of Gerard Grisey or Marco Stroppa and which is today related to the research on automatic orchestration (Carpentier, Tardieu, Esling) such as used by Jonathan Harvey. In these approaches, information contained in an audio signal (chords or timbre) is used to generate a score, which is then played by an orchestra. Another kind of approach relates to the learning of language models (LZ, Factor Oracle) to represent a performer style. Example of automatic improvisation by a computer using this model trained on Coltrane solo or trombone-solo (using the O-Max system) were demonstrated. As said, it was the first time a computer was applauded on stage.

According to Norbert Schnell, nowadays, music creation involves less and less the body; while music listening tends to involve more and more it through interaction. This interaction with music is performed by "re-liquifying" the audio. Many examples of this re-liquifying process of audio were given: - using gestural control to perform direction of pre-recorded music (collaborations with atelier feuillantine), - controlling video beat-boxing segments using a gambling machine (mind-box project), controlling synthesis using a ball within a game (Urban musical game). Many of the paradigms used here for creation are based on "breaking and remaking time". Two possibilities are derived from it: "respect and readapt" or "eliminate and recreate from scratch".

Diemo Schwarz illustrates this paradigm of "eliminate and recreate from scratch" with the CATaRT system. In this system sounds, segmentation and audio descriptors are used to inform a real-time concatenative synthesis system allowing re-creation of any type of sounds or texture from pre-recorded audio or live-recording of audio in real-time.

Philippe Manoury, one of the main composers of today, gave its view from the creation side. He first emphasizes on the difference between composition and improvisation, which involves different time-scale (long and short-term). He then emphasizes on the importance of acquiring knowledge of the sound. Since sounds are always in evolution (from Rameau, Webern to Grisey) audio descriptors are particularly important for a better understanding of sounds. He then provided examples of the use of music information in its pieces: score-following in "Partita-1" and audio-descriptors in "Tensio". He finished by recommending the development of audio descriptors closer to the way we ear and create music.

Tristan Jehan (EchoNest) presented the EchoNest platform, which is possibly the largest content-based M.I.R. recommendation system; used today by record-labels (EMI) and online-music-providers (Spotify). The EchoNest platform is organized around a set of APIs some of which are public. Through the use of these APIs, he illustrates how to make creation with audio-features computed on the millions of songs EchoNest have, using their Remix Python API: wubmachine, videolizer, the swing thing, the walzify, the global sampler. Some of these creative applications are developed during the Music Hack Days, which appears as a very productive place for M.I.R. based creation.

Francois Pachet (Sony-CSL) presented works related to style. For him music is texture and structure. Study of style has been performed in the MIROR project where Constrained Markov Model are used to create for example improvisations or generate blues grid based on Parker chords. In a recent project, Virtual Bands are created. This is done by analyzing musicians separately, extracting features and then modeling them using CMM. The system can then be used in real-time to accompany a musician while following its intentions.

Works of the MTG in M.I.R. started with what Sergi Jorda name "traditional M.I.R." (Canoris API, BMAT spin-off or Song-Explorer using the reaktable to navigate over a music collection) dealing more with consuming than producing. According to him, we can distinguish: - micro-MIR for Music Creation (such as plunderphonics, mashup, concatenative synthesis, musaicing, beat-mash/ loop-mash of Steinberg), - meso-MIR for Music Creation (freesound and its proposed taxonomy of sound, freesound radio and the possibility to create graph connection between samples, sound-scape modeling) - macro-MIR for Music Creation (expert system such as the LaDiDa app, auto-tune for the news). He then illustrated how M.I.R. and creation does not only concerns music creation. The "Magic Fountain of Montjuic" uses a system developed by the MTG. In this, content-description (audio features extracted from music tracks) drives the light-choreography of the fountain. To conclude, he stresses that "playing music is not the same as playing with music" and that applications should “keep things as simple as simple but not simpler”.

2.4.8 Panel discussion: ***MAKE, PLAY, SHARE: The future of MusicTech*** *at the Sonar festival in Barcelona, Spain (14 June 2012)*

The panel intended to look into the future of the Music Industry from three perspectives, Search & Discovery, Creativity and Rights Management, and gather first-hand feedback from relevant industry players. The following topics were addressed:

- Music Recommendation: Is it as good as Human Recommendation?

For some speakers, Music Recommendation is as good as it can be. For some others, there is still a long road ahead for every user to have their own John Peel. Also, Music Recommendation still lacks to take the time factor; the sociocultural evolution of what a given song/artist represent. This is a big topic which should be addressed.

- Music Charts

They are still relevant. Music is emotion and connects people. Listening to the same songs as your peers is a way for creating bonds with them. The charts are now more granular per tag/microstyles/community of friends.

- Geographical Tagging: very relevant for music personalization.
- Music Services Interoperability

Absolutely necessary in order to create new user experiences. Musicbrainz is trying to reach that by building a unified music metadata database. New tools have appeared in the last years to resolve content regardless across music repositories, streaming services and geographic territories (http://www.tomahawk-player.org/).

- Playlist Generation (Music Listening Experience)

Since music is basically emotion, new music experience tools must focus on emotional coherent experiences.

- User Experience: “why music services still look like music spreadsheets?”

More research is required in tangible and visual interfaces. Songs are no longer than 3-minute mp3 files but they can be hours long, divided into tracks available through the artist mobile app.

- Unified Music Database: having a Unified universal metadata database for musical assets is key for building the music rights tracking systems of the future.

2.4.9 ***Music Hack Day****, Sonar 2012 festival in Barcelona, Spain (15-16 June 2012)*

Music Hack Day has been a great way to demonstrate the creativity around music that comes from the tech community, fostering cross-platform and cross-device innovation. The past three years have seen more than 20 Music Hack Day events taking place around the world. Starting in London, it has spread across the world to Berlin,

Amsterdam, Boston, Stockholm, San Francisco, Barcelona, New York, Sydney and Montreal. MHD is a very exciting event, watching over 2000 participants taking part, building hundreds of hacks and with over 125 music and tech companies' supporting the events.

The Music Technology Group has organized Music Hack Day events in Barcelona since 2010. This edition of Music Hack Day at Barcelona is being organized in collaboration with Sonar Festival like last year, taking advantage of the great exposure presented by this Festival (more than 79,000 spectators, 833 companies and 1,018 journalists accredited from 515 media last year 2011).

This Barcelona Music Hack Day edition involved 30 people representing 18 companies or entities coming from the music technology field plus 90 hackers and 10 artists who were in charge of creating applications using involved entities' resources.

The session started with the presentations of the tools available for hackers given by 18 companies' representatives. Then the conceptualization stage started, where participants (or hackers) suggested hack ideas and built teams, based on individual interests and skills. At 14pm the main work of the hacking session began, which lasted 24 hours until the 14pm of the following day (overnight hacking included). At the end of the hacking session, there was a series of demonstrations of the 37 hacks developed as main outcome of the session (all presentations available here http://www.youtube.com/playlist?list=PL76079818BE31A2D1). The best in-show hacks were awarded by a panel of judges from involved companies who selected the winning teams, and the prizes were given.

2.4.10 Panel discussion: ***The Future of Music Information Research*** *at the CMMR conference in London, UK (21 June 2012)*

On 19th-22nd of June 2012, Queen Mary University of London organized the 9th International Symposium on Computer Music Modeling and Retrieval (CMMR 2012) at the QMUL Mile End Campus, London, UK. The main conference topic was "Music and Emotions". In the third day of the symposium, a panel discussion was organized on the "Future of Music Information reSearch", moderated by Prof. Geraint Wiggins (QMUL), with invited speakers being Prof. Joydeep Bhattacharya (Goldsmiths, University of London), Prof. Tim Crawford (Goldsmiths, University of London), Dr. Alan Marsden (LICA, Lancaster University), and Prof. John Sloboda (Royal Holloway University, Guildhall School of Music & Drama). Panelists were asked to prepare a 10-minute introduction, followed by cross-questioning and audience participation. The panel discussion lasted about 90 minutes, was attended by about 50 people, and can be viewed online at http://www.mires.cc/videos/CMMR2012.

The themes and challenges that arose from the discussion can be grouped into the following categories:

Musical emotion

- Music-induced mood and emotion: are they the same or different?
- Are musical emotions "real"?
- Perceived and felt emotion - are they the same or different?

- Are musical emotions similar to other emotions?
- What are the best models of musical emotion?
- Is musical emotion equivalent to aesthetic experience?
- What is the role of expertise, gender, culture and social context?
- Most studies are relaying on self-reports of emotions; are they trustful?

MIR tools/systems for historical musicology

- Are current industry-oriented applications suitable for musicologists?
- What is our main source for text/label information retrieval?

Provocations about musical value and culture

- Are we confusing between being emotional about music and recognizing emotions in music?
- If we care about music so much, why do we not treat it like a precious item?
- Do we have an effect on music that we research by doing our own research?
- What is MIR research doing to musical culture, what are the consequences?

Expanding the musical object

- What is the definition of a musical object?
- Can we expand that definition to include information on the performance in context?
- Can context and liveliness be retrieved by MIR?
- Is MIR “finding the music you want”?
- If so, can MIR extend its remit from past recorded performances to future live ones?

Addressing human benefit

- How does a researcher decide what research to do?
- How does a research community decide what research to support/promote?
- Are there any policies and procedures for MIR?

*2.4.11 Panel discussion at the **2nd CompMusic Workshop**, at the Bahçeşehir Üniversitesi, Beşiktaş, Istanbul, Turkey (12-13 July 2012)*

The workshop included 28 talks, most of them on research carried out within CompMusic but we also had presentations by external researchers working on topics of relevance to the project. Among the topics that were emphasized are:

Tonic detection and Intonation characterization in Indian music: In Indian music the tonic functions as the reference pitch from which all the harmonic and melodic relationships are established during a performance. Its automatic detection is a necessary step for analyzing many musical aspects and thus it was the first problem that we wanted to solve. I consider the tonic detection results obtained and reported in these articles sufficient for the purposes of the project. A natural next step is the characterization of the intonation used in a performance. This is a not solved problem but the initial results reported at the workshop are quite promising. The proposed representation of intonation in raga music can be a valuable element used to describe various musical aspects.

Melody issues: There were many talks related to the characterization of melodies, both in Indian and Turkish musics. An advantage of Turkish music is the existence of scores, making it very valuable for the study of some melodic aspects. For using the audio signals an important first step is the extraction of the pitch contour of the lead instrument, which requires a multipitch methodologies. Both makam and raga music have particular ways in which melodies are constructed and thus they require specific approximations in order to characterize them. Most of the presentations reported preliminary work and there is a long way to go before we have proper parameterizations with which to identify musically meaningful melodic elements. Two of the papers by external researcher focused on the synthesis of melodies, which is an interesting aspect not covered in the project.

Rhythm issues: Both Turkish and Indian rhythms are very particular. The rhythm in Turkish music revolves around the concept of usul and the one in Indian music is based on the concept of tala. These concepts have been quite well formalized by musicologists but there is practically no computational work on them. Two of the articles presented very preliminary computational work, but there is a lot to be done before the existing musicological formalization can be turned into algorithmic approaches to describe the rhythm of both Turkish and Indian musics. One article presented some work on perceptual aspects of rhythm, which we are not covering within the project.

Community characterization: A basic premise in CompMusic is the need to study and characterize the communities, people, that support the music that is being studied. We want to analyze on-line communities and develop methodologies to extract musically relevant knowledge from them. Freesound is an on-line community of people that share and talk about sounds that we have used to start developing methodological approaches of relevance to the project. The first musically relevant on-line community that we have studied has been Rasikas.org, using it to extract musically relevant information for the case of Carnatic music.

Tools for browsing and discovering music: In CompMusic we are very interested in developing tools and methodologies to be used in practical music applications, especially applications for browsing audio music collections. A relevant functionality is to be able to discover patterns and musical relationships from which a listener can learn to appreciate better the music of a given culture. We presented a paper on the initial application prototype that is developed and other researchers presented systems that are being used in tasks of relevance to CompMusic.

2.4.12 ***Music Tech Talks****, as part of the Cultural Olympiad in London, UK (26 July 2012)*

Upon invitation by the London Cultural Olympiad, and in order to capitalise on the momentum generated by the Music Tech Fest, a follow-up event to of Music Tech Talks was organised on July 26th 2012. The event involved guest speakers Matthew Hawn from Last.fm and Andrew Shoben from Greyworld, as well as a guest appearance by composer and artist Tim Exile, who handed the awards for the winning hacks at the Music Tech Fest 3D Music Hack Camp.

Andrew Shoben, founder of Greyworld discussed the potential of small sonic interventions embedded into the urban fabric of a public space, to allow some form of self-expression in areas of the city that people see everyday but normally exclude and ignore. He presented a series of examples of work, both physical and aural, that attempt to establish special intimate zones, to ‘short circuit’ both the environmental and social expectations supplied by the surrounding urban realm.

Matthew Hawn of Last.fm delivered a talk about the new applications that Last.fm are building at to enhance the music playlist user experience, and about the collaborations with students and academic researchers in Aarhus, Denmark, to create novel forms of physical interaction using the Last.fm API.

The recordings of the talks have been added to the Music Tech Fest TED-style channel on YouTube, which has become an important MIReS repository of documented events and public information, as well as a platform for attracting proposals for further contributions. The film release programme has enabled the platform to continue to build a global following and to reach new territories like Brazil, that had not originally tuned in to the live broadcast during the festival. The partners plan to continue to incentivise the Music Tech community thus created throughout the duration of the MIReS project, and encourage continued use of this framework beyond the duration of the project.

2.4.13 ***MIRrors: the past of Music Information Research reflects on its future****, at the ISMIR 2013 Conference, Porto, Portugal (8-12 October 2012)*

A special session was organized in ISMIR 2012 to reflect on past practices and existing knowledge that could have been overlooked or that could help guide the consolidation of the MIR discipline. A special call for papers and a review track was held separately from the regular papers resulting in 6 out of 8 submissions being selected for the session. The papers were downloadable for weeks before the conference in order to facilitate reading by interested researchers in advance of the conference. During the session, authors were given 8 minutes to summarize their positions and then the audience could interact with them, asking for clarifications, debating or commenting on their issues. Interaction between authors with overlapping topics was also promoted. A special Google vote web was open to help researchers contribute to the discussion after the session, where 25 questions were addressed to the authors.

Insights included reviews of the shortcomings of the existing approaches to music transcription, MIR generic systems and user modelling; highlighting of capital misunderstandings between researchers of potentially synergetic disciplines (MIR and music cognition); lack of infrastructures to reuse, recycle and take advantage of the existing MIR knowledge that exists as computer code.

2.4.14 ***Evaluation Session****, at the ISMIR 2013 Conference, Porto, Portugal (12 October 2012)*

The goal of this panel was to discuss on current evaluation practices within the MIR field and compare those to the one currently used in Information Research. The panel was composed of representatives from main MIR evaluation initiatives (MIREX, MillionSongContest, MusClef) and a representative of IR evaluation: Gareth Jones, one

of the founders of MediaEval. Panellists included Gareth Jones (Dublin City Unviersity), Brian McFee (University of California at San Diego), Nicola Orio (University of Padova), Julian Urbano (University Carlos III of Madrid), J. Stephen Downie (Univerity of Illinois at Urbana-Champaign). Topics included:

Definition of the tasks to be evaluated

What methodology should be used to define the task (bottom-up vs. top-down)? For which purpose should a task be evaluated: low-level tasks (functionality-oriented such as beat, chords) vs. full-system tasks (use-case-oriented such as music recommendation systems). Specific tasks that are part of large-scale international evaluations define de facto the specific topics that new contributors to the MIR field will work on. The methodology followed to define tasks is therefore of utmost importance.

Evaluation

How should a specific task be evaluated? Which data, which measures, what is the reliability of the results obtained?

Data

How to get more data? How to deal with data availability (not only music collections, but also raw system outputs, judgments, annotations)? Should we go to low-cost evaluation methodology (see TREC Million Query Track 2007, 2008 and 2009)? Currently most MIR systems are concerned with audio-only or symbolic-only scenario. Multi-modal systems (such as aggregating information from the audio-content, from lyrics content or web mining) should also impact on the final user application of each technology.

Methodology

What is the best methodology to drive improvements? What kind of evaluation framework (open vs close evaluation)? What could be improved in previous evaluation initiatives? How can we make results reproducible? How can we make MIR evaluation sustainable along time?

The most important contributions included:

- how to solve the close model used by MIREX (no data publication) forced by the copyright nature of the audio tracks.
- the centralized organization of MIREX. Media-Eval proposed a completely de-localised organisation with different leaders for the various tasks.

- the choice of tasks to be evaluated: bottom-up (algorithm oriented) versus top-down (use-case oriented).

2.4.15 ***Demos and Late-breaking sessions: Music Information Research Challenges****, at the ISMIR 2013 Conference, Porto, Portugal (12 October 2012)*

The ISMIR late-breaking session “Music Information Research Challenges” was held in response to the following question: what are the most important challenges facing the MIR community in the coming years?

In the MIReS project the partners had begun to identify several areas for future investigation by considering technical, as well as social and exploitation aspects of MIR research. Amongst the many topics were musically-relevant data, knowledge-driven methodologies, interface and interaction aspects, evaluation of research results, social aspects, culture specificity, industrial, artistic, and educational applications. The list of challenges was a work in progress and could be found on the MIReS wiki.

Researchers who may have already begun to address some of these new challenges were invited to add comments or additional challenges and highlight the topics which deserve further discussion. During the discussion it became apparent that the exploitation of MIR research results was a more recent addition to the research community and of substantial interest to researchers. The discussion stimulated further initiatives by the MIReS consortium which would ensure that academic research translates the feedback from industry into well-formulated research challenges for the MIReS research community.

2.4.16 ***Music Tech Talks II****, at Queen Mary, University of London, UK (15 November 2012)*

With Music Tech Talks II, the partners continued to incentivise the Music Tech community, with contributions by Robert Kaye of MusicBrainz, Prof. Mark Plumbley and Dr George Fazekas of partner QMUL.

Robert Kaye is founder and lead developer of MusicBrainz, a leading open source project for music metadata on the internet, and President and Executive Director of the MetaBrainz Foundation. MusicBrainz utilises an approach similar to Wikipedia to curate high quality metadata and to assign unique identifiers for metadata entities. These identifiers allow for unambiguous communication about music through its global metadata delivery network that includes customers like Google, BBC, Last.fm, Grooveshark, Amazon and AOL.

Professor Mark Plumbley introduced the work of the Centre for Digital Music (C4DM), Queen Mary, University of London. The Centre is a world-leading multidisciplinary research group in the field of Music & Audio Technology, investigating topics such as music information retrieval, music scene analysis, semantic audio processing, object-based audio coding, human machine interaction and digital performance. With its broad range of skills and a strong focus on making innovation usable, the Centre for Digital Music is ideally placed to work with industry leaders in forging new business models for the music industry.

Dr George Fazekas delivered a talk about technologies developed at the Centre for Digital Music, including knowledge transfer initiatives, such as their content-based recommendation technology currently trialled by the BBC and iLikeMusic, the BBC Desktop Jukebox and research tools such as the Sonic Visualiser, which is widely used within the music information retrieval community.

2.5 Framework for a research-to-industry network

The consortium proposed to build a framework for the establishment of a virtual network of MIR excellence connecting music researchers and theorists with industry representatives, strategists and policy makers, ensuring efficient communication and transfer of innovation to industry in the long-term, and thus continuing to inform policies beyond the lifetime of the proposal to ensure future academic excellence and EU competitiveness in the world music market. The consortium had sought to make preparatory steps, initially drawing upon dissemination events from WP4 and community co-creativity events from WP5, and working actively towards expansion by forging links throughout the music technology stakeholders community.

Following the success of MIReS coordination initiatives such as the Music Tech Fest, the Synaesthesia Creative Workshops, 3D Music Hack Camp, the Make Play Share Sonar Panel and the Music Hack Day, regular communication has been established between academic researchers and the music and creative industries enabling more efficient transfer of research-to-industry.

Recommendations have been made in view of ensuring that academic research translates the feedback from industry into well-formulated research challenges. Aside from discussing and building on the industry feedback gathered from specially-organized MIReS events, the partners have researched into methods in use by other academic research communities and have identified examples of good practice for involving industry and wider stakeholder networks with academic research. As a result of these investigations, dialogue between industry and academia is underway to identify “Grand Challenges for MIReS” for a call for submissions, planned to be evaluated in 2014.

Long-term impact has been achieved by the recognition and allegiance by stakeholders to the Music Tech Fest community, with continuing contributions and proposals for future events. The MIReS organising partners were invited to organise two follow-ups to the festival. The first, an evening of Music Tech Talks, was held on July 26th 2012 as part of the Cultural Olympiad during the London Olympics. A further event was requested by the community and organised on the 15th of November 2012 with speakers including some of the more prominent figures of the music technology community.

Framework building has continued with the growing Music Tech Fest online media channel which is receiving ongoing requests for contributions from the music technology community. At the time of writing this report, the Music Tech Fest channel has registered 6,441 views. Related social media channels (e.g. @musictechfest on

Twitter and the Music Tech Fest Facebook page) have over 1,000 subscribers and over 500 regular visitors.

The partners have now secured funding for the Music Tech Fest to continue beyond the lifetime of the MIReS project. The Music Tech Fest 2013 will be held in London on the 17-19 May, funded by the European Regional Development Fund. New contributions are planned to populate the Music Tech Fest channel and further contribute to the MIReS knowledge base.

Satellite events have been requested by stakeholders from New York, Portland (Oregon), Ireland, Brazil and Croatia. The organisation of the Music Tech Fest satellite event in Croatia is planned for 19-21 July 2013, after the entry of Croatia into the EU on the 1st of July 2013.

The organizing partners have been actively pursuing funding opportunities to extent the MIReS stakeholder community within and beyond the EU, and are currently in talks with major brands to make the Music Tech Fest a global event. For this purpose Music Tech Fest branding has been registered with the UK Intellectual Property Office as a Registered Trade Mark, and the business has been incorporated as a UK Limited Company with company number 08293318. The Roadmap activities have thus instigated the establishment of a business entity, and secured long-term involvement of music information research stakeholders from both academia and industry.

3 CONCLUSION

The Roadmap for Music Information ReSearch has achieved, and in some cases exceeded, the planned project impact and objectives. The roadmapping process has substantially addressed the state-of-the-art of the field, research and industry standards, and the future challenges definition; it has organised high-profile activities attracting contributions from a variety of experts, addressing different topics relevant to the roadmapping process; it has aligned itself to other ongoing initiatives such as CHORUS+ and the Media Search Cluster with a view of influencing policy making and research planning; it has successfully established a regular exchange of communication within the research-to-industry network with a view of enabling a long-term transfer of innovation to industry and ensuring future academic excellence and EU competitiveness in the world music market.

Notable achievements have included a substantial increase of the wider network of MIR stakeholders; successful dissemination and spread of knowledge about this field of research to industry stakeholders and other academic fields; introduction of valuable research challenges to the MIR academic community; and long-term impact through ongoing activities and events established by the MIReS project.