Musicology: Bringing Personal Music into Shared Spaces

Anonymized for Blind Review

The Musicology project examines the user, ecosystem, and the technology aspects necessary to enable music sharing in public spaces. Advanced cell-phone platforms, for example, can serve as conduits for expressing a users personal music in shared spaces such as the digital home or public cafés. User surveys, field interviews, and a technology prototype inform a discussion about a variety of topics, uncovering concerns relating to privacy perspectives, digital rights management, and power consumption (among others). Sharing personal music in public spaces is an exciting area that has already started to see some intermediate solutions work their way into social practice, but a few key hurdles still need to be resolved in order to bring such systems into mainstream usage.

Introduction

Emerging mobile devices, such as multimedia cell-phones and high-capacity mp3 players, are poised to foster new types of multi-user music interactions in shared spaces. Such interactions have the potential to enable new ways for individuals to express their personal identity and relate to the environment around them. For example, a person could play a selection of their music, stored on their mobile phone, through a coffee shop's wall-mounted speakers, allowing them not only to enjoy their personal music without isolating themselves using headphones, but also to hear music provided by other patrons in the space. Music is a highly dynamic and popular medium for personal expression, recently accounting for significant growth in the mobile device market. The Musicology project explores the underlying user perspectives, enabling ecosystem, and supporting technologies necessary to catalyze the social consumption of personal music.

Music has always been a very popular application for mobile technologies, ever since the inception of the portable tape player, and it is a powerful medium for expressing personal identity and understanding the identity of others [3][7] or fostering a sense of shared identity among a group of people [5]. This social aspect of music creates a great potential for supporting technologies: Cell-phone ring tones are a prime example, enabling people to aurally present an aspect of their identity while providing a steady revenue stream for wireless carriers [16]. Despite these basic social drivers, the long standing tension between individuals and the music industry regarding the sharing of music [14] highlight the complex dynamics of this space, motivating the need for well-crafted solutions addressing the concerns of all the relevant stakeholders.

Music in shared spaces, such as a public café or the home environment, is often used to set the mood or ambiance [6], and can sometimes even become a focal point of the space itself. Typically, the owners or workers of a public place define the music; and in some circumstances, people are able to directly affect the music by interacting with the people responsible for playing the music, or by directly interacting with a situated device in the space itself, such as a traditional music jukebox. Deterred by the social and technological hurdles, people only rarely bring their own music into a public place, either in the form of CDs or by plugging their own personal music device into the infrastructure. More often, when people want to listen to their own music, they end up use headphones, cutting themselves off from others nearby, albeit sometimes intentionally [4]. Even in the home environment, it's still tedious for people to play music for each other because of the logistics involved.

The capability for mobile devices to support the contribution of personal music to shared spaces is born out of advances in mobile storage, processing, and communication technologies. Music-centric devices, such as the Apple iPod, provide a first-class personal music experiences but lack the wireless communication capabilities necessary to enable seamless integration with shared spaces. Cell-phones such as the Motorola Rokr and Nokia N91, among others, possess all the necessary components for supporting interaction in shared spaces. The Rokr further combines local music capability with the popular iTunes software package, connecting the device to the greater music ecosystem [10]. However, although the underlying technology of these platforms is capable of wirelessly streaming music from the device, the current implementations of these devices do not support effective use in shared spaces, motivating the need for solutions that specifically target these situations.

Independently of the cell-phone industry, the digital music and digital home industries have been developing and indirectly supporting the multi-user playback scenario. The digital music ecosystem, spurred by the high-density mp3 encoding format and on-line music stores, has created an environment where large personal digital music libraries are commonplace. Similarly, digital home initiatives such as the Digital Living Network Alliance (DLNA) [11] have been producing technologies for consuming digital media in highly networked environments. Nominally used to support pervasive media consumption throughout the home, these digital home technologies provide an attractive stage for integrating mobile devices carrying digital content into the home. The Apple AirPort Express [9] is a prime example of these two ecosystems coming together: it is a device that allows you to play music from your laptop on a home stereo system over a WiFi connection. Although it allows one to wirelessly play music from a mobile platform, the setup neither supports switching between multiple music sources nor allows interactions with the more personal cell-phone form-factors, two components necessary for a wide-spread shared space experience. Streaming music from a laptop is the musical equivalent of using VoIP to make a quick phone call from a laptop: it works, but is not likely to displace the cell-phone.

The individual elements described in the proceeding paragraphs are not in themselves sufficient enablers of a multi-user music system since they are all aligned towards related, but slightly different, needs: a deficiency directly address by the Musicology project. This paper combines a user survey, semi-structured interviews, and a technology prototype to form a coherent framework for enabling music systems in shared spaces.

The specific research contributions of this work cover three areas:

1. An understanding of the role of privacy and familiarity, highlighting identity and personal relationships as key components.

2. A delineation of the four stakeholders for shared-space music ecosystem, highlighting their concerns and contributions to a sharing cycle.
3. A prototype music sharing implementation that has improved power and control properties when compared with existing solutions.

The rest of this paper is organized into six sections. Immediately after this introduction is a summary of related work. Then, three sections then cover various aspects of the user experience, overall ecosystem, and underlying technology of public music systems. The paper ends with a discussion about overcoming potential barriers to adoption and a brief conclusion.

Related Work

A seminal work on music in shared spaces is Music FX [15], which allows people to register their music preferences within a gym environment. The system automatically adjusts playback in the space to match the preferences of those present after they swipe their gym badge for entry. Music FX focuses on the collaborative filtering nature of music preferences and explored how they can be used to (positively) shape music experiences in the shared space. Musicology shares many goals with MusicFX, extending the concept by allowing people to contribute personal content and interact with the system using a personal mobile device instead of an identification badge.

The Jukola [17] system presents an implementation, deployment, and in-depth analysis of a system for music interaction in a shared public space. A kiosk interface is used to nominate songs drawn from the establishment's library of music, while PDAs, which can be checked out for temporary use, are used to vote for the next song to play. Both Jukola and Musicology promote public interaction with music in a shared space using a combination of mobile and fixed devices. Jukola, however, dictates which interfaces can be used for each task and does not consider the use of *personally owned* devices for interaction in the space. Only the local music library of the establishment, augmented with songs submitted over the web, is available to patrons for playback, making it more difficult for people to express their personal identity through the system, although affording the establishment much more control over what can be played. The key differences explored by Musicology are the affordances offered by personal mobile devices, a more structured framework of the needs of the various stakeholders, and a technical analysis of the supporting mobile technologies.

The tunA system [2] investigates using personal mobile to create personal local-area radio stations. So, for example, if several people are listening to music on a bus, they can "tune in" to one another's currently playing music. tunA only partially enables a shared space by allowing people to listen to the same content while still remaining isolated through the use of headphones. Furthermore, tunA is based on WiFI technology and, although it *functionally* allows music to be streamed wirelessly from the device, the system admittedly places an excessive power drain on the device (limiting battery life to around three hours). Musicology moves the consumption of music into a physically shared auditory space, allowing people to affect music playback in an environment, instead of simply importing others' content into their personal audio space. Additionally, Musicology fully considers the power impact of the supporting wireless technologies.

The iTunes music player can be used to share music over a local area network, highlighting how personal music and laptops can create complex social interactions [23]. iTunes enables users to publish their playlists, from which other people can play the listed music from their respective laptops. Users are not necessarily co-located when sharing (rather, they are just on the same subnet). Overall, the interaction is fairly heavyweight, since full-featured laptops must be used for the sharing (needed to run iTunes), but the base functionality captures many of the social and legal ramifications of personal music sharing. For example, a user is able to remotely access, explore, and listen to entire music collections stored on another user's system, but they are not able to copy music from the collection to their local machine. Similar to the previously mentioned AirportExpress adaptor, these technologies start to move in the direction proffered by the Musicology model but still do not seamlessly enable a shared space since users are still *consuming* the music individually. One way to frame Musicology is as an extension of these technologies and practices to personal handheld devices and shared spaces.

iRadio [8] is a forthcoming system from Motorola that positions a mobile phone as a source for a user's personal radio station, providing a seamless music stream that can be accessed at home, in the car, or on the go. Content is organized into channels, conceptually like radio stations, allowing for easy integration with existing car stereo systems, which typically have buttons for selecting between preset radio stations or CDs in a CD changer. One of the iRadio channels can be a collection music drawn from the user's personal collection. Musicology and iRadio are very complimentary, as both center the music experience around wireless access to personal media provided from a personal device. However, iRadio does not directly support multiple devices in a shared space nor does it allow full access to the content stored on the device itself, since access is aligned along the channel metaphor. The Musicology investigation provides an analysis of some of the underlying standards used to implement iRadio, alternate design solutions, and an exploration of how a system like iRadio might serve to enable music consumption in shared spaces.

Musicology builds on the Personal Server model [24], a system that advocates access to the content stored on a mobile device through interaction elements in the infrastructure, by treating a user's mobile device as a personal (music) server. This perspective follows the original Ubicomp vision [25] by investigating the steps necessary in order to create systems of well-connected devices. Similarly, it extends the line of investigation considering how a highly personal technology, such as the cell-phone, is part of a complex and dynamic ecosystem [19]. With both these cases, Musicology applies the broader concept on a specific application domain (music in shared spaces), to provide a more focused understanding.

User Perspectives: Privacy and Familiarity

A user survey was used to investigate the relationships between public music consumption, personal privacy, and familiarity within a public space. The survey included a series of questions with 7-point Likert scale response options, interspersed with open-ended questions at the end of each subsection. The first half covered a sequence of general issues about the participants general music listening behavior, the specific locale, privacy preferences, and projections about playing personal music at that location. The second half investigated a series of specific behaviors (e.g., playing one of their songs, voting on the currently playing song, etc…), across a variety of interfaces (talking to the store staff, using a kiosk, from their mobile device, etc…). By design, the survey did *not* focus on any interface usability aspects of the system, and therefore only screen-shots of various interfaces were shown (no technology nor partially-functional prototypes were used to demonstrate or explain the concepts).

We administered the survey to 20 participants over a period of two weeks in a popular cafe in the (Anonymized) metro area, near a college campus. Respondents were asked to express their preferences in "this café" at "this time" in order to strengthen the contextual nature of the questions. The cafe has a full bar, features local artist's work on a rotating basis, and continually plays music from a laptop placed behind the bar (controlled by the staff). Surveys were distributed to every patron in the cafe at a variety times on several different days, and participants were asked to fill out the survey at their leisure while in the cafe. Each participant filled out the survey where they were already seated, while the experimenter remained distant but visible in case the participant had any questions.

The selection of key results presented in the following sections focus on various actions, interfaces techniques, personal privacy, and familiarity with the space. The survey covered a broad range of topics, and some results have been omitted because they were not statistically significant. Statistical significance is measured using the Student's t-test, and p-values are indicated where appropriate. Typical values of p are [$p < .10$] or [$p < .01$], which indicate a less than 10% or 1% chance of error in the result, respectively.

Actions and Interfaces

The survey asked participants the likelihood of using each interface for a series of actions, described in Figure 1. There was no clear advantage averaged across ant of the actions; however, using the kiosk or an anonymous mode of mobile interaction were viewed more favorably for playing a song (A4) [$p < .10$], while other interfaces were not as compelling – indicating the potential for a mobile interface, but potential difficulty for systems requiring registration. The main reason cited for this by the respondents was "laziness" (people were worried that the registration process might be too difficult, especially from a mobile phone device with limited input capabilities) and not wanting to share any personal information (although this conflicts somewhat with the results in the next section). The choice category, which represents the most favorable interface for each individual participant, highlights the effectiveness of allowing people to chose their preferred interface: providing a choice significantly increases the score for all actions. For example, it would be perfectly reasonable to allow people to either use the kiosk or their mobile device, depending upon where they were, what they wanted to do, and whom they were with.

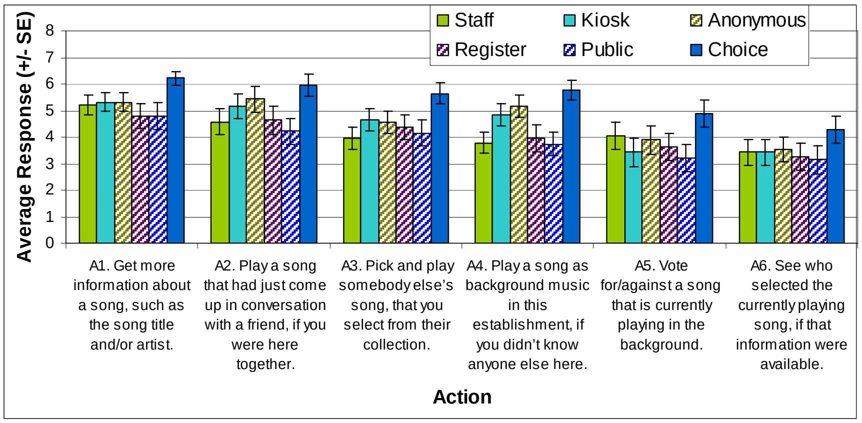

Interface	**Medium**	**Registration**	**Physical Motion**
Staff	Personal interaction	None required	Go up to store staff
Kiosk	Situated display	None required	Go up to situated kiosk
Anonymous	Mobile device	None required	None
Register	Mobile device	Visible to store staff	None
Public	Mobile device	Visible to everybody	None
Choice	*Composite score representing the maximum response across all interfaces*		

Fig. 1. Survey results for a selection of actions across the various interfaces. Actions are sorted left to right in order of decreasing average response. Error bars show +/- one Standard Error for each column.

Based on the average responses across all interfaces, playing a song when not with a friend (A4), ranked lower than playing a song when with a friend (A2) [$p < .01$], indicating the influence of social interaction, or social support, on the system. Furthermore, this discrepancy disappears when considering the kiosk or anonymous mobile interface: people were more willing to overcome the registration hurdle in the interfaces to play a song for somebody else, but less likely to do so when by themselves.

Voting for the currently playing song (A5) scored less than simply playing a song (A4) [$p < .10$], which is interesting considering the context of the Jukola system, which successfully made voting a key part of their music system. One significant difference between the systems is cultural: the UK (Jukola) vs. US (Musicology), or just the particular café environments. Furthermore, voting in Jukola is used to select which song to play *next*, while the Musicology survey described voting as a means to express an opinion about an already playing song – possibly deterring people from voting lest they offend somebody. Interacting with the store staff was generally viewed as less likely, *except* in the case of voting, where it was more on-par with other techniques. One reason for this discrepancy may be an increased belief in the *authority* placed in the store staff (they are responsible for the space), or buffering against potentially offending anybody (using the staff as an intermediary to soften the blow).

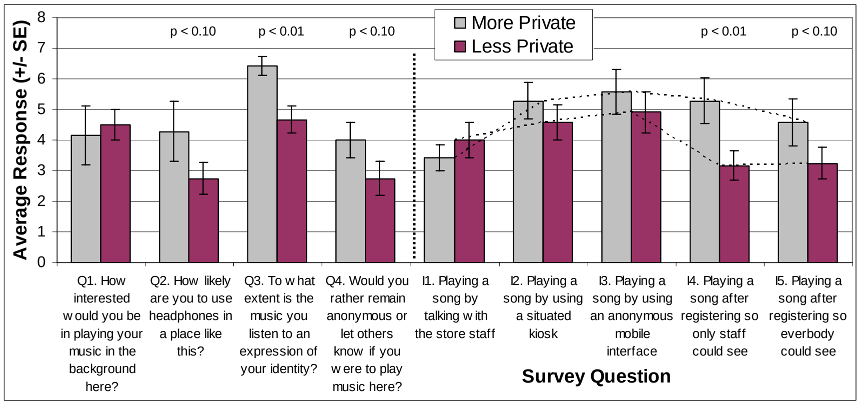

Fig. 2. Effect of privacy disposition on survey responses. Answers are divided into two categories based on their response (above or below average) to a group of privacy related questions. Error bars indicate +/- one Standard Error.

Privacy and Identity Boundaries

A series of responses, shown in Figure 2, highlights the impact of privacy based on responses to a group of questions inspired by a standard privacy classification test [20]. The basic privacy indicators consist of three standard questions about a user's view about an employer's control of employee information, an organization's control of information they collect from their employees, and laws' effectiveness to protect employee privacy. These questions are combined into one composite score, which is then used to divide the survey population into two halves, allowing a differential comparison between "more private" and "less private" dispositions. For example, although not significantly impacting their interest in playing their music (Q1), private people indicated they were more likely to use headphones (Q2).

Privacy disposition has a strong impact on the measure of a participant's identification with music (Q2). This is not necessarily a causal relationship, but it sheds some light on other the comparisons. Essentially, a strong awareness of identity encourages a stronger sense of privacy needed to protect that identity. For example, many of the emerging laws used to combat *identity* theft are oriented towards protecting a consumer's *privacy* – highlighting the direct relationship between the two notions. This emphasis on identity is different than concerns raised by many other personal device systems, which typically focus on location-aware capabilities (e.g., [1][22]), and in turn have very different privacy implications due to the nature of the shared information (you don't need to share your location with somebody in the same room).

More private participants indicated they were *less* interested in remaining anonymous when playing music (Q4), which is somewhat counterintuitive because they would have to give up some of their privacy to identify themselves when playing music; however, this result is congruent with the notion that people are interested in expressing their *identity* through music, consistent with the concept of privacy as an

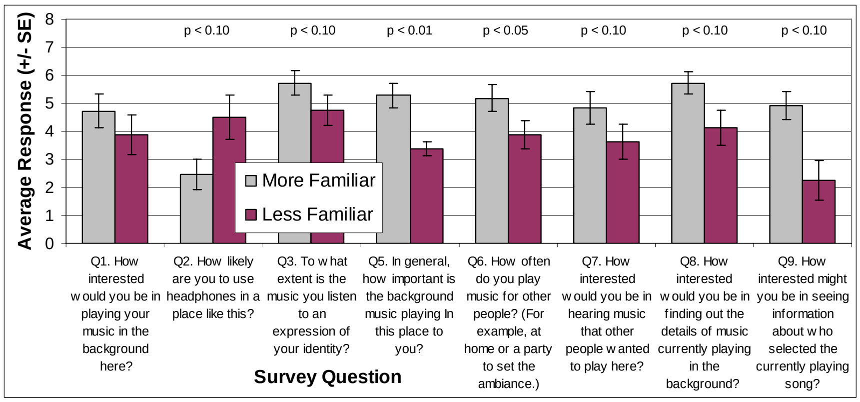

Fig. 3. Effect of familiarity on survey responses. Answers are divided into two categories based on their response (above or below average) to a question about how familiar they were with the particular environment. Error bars indicate +/- one Standard Error.

identity boundary [18]. From a design standpoint, this means designing around identity management, not just protecting privacy. Privacy had a significant positive impact for interfaces that required registering (I4, I5) – again supporting the perspective that privacy and identity are tightly coupled. People still preferred the anonymous interfaces, i.e., they would prefer not to register, but they were less opposed to having to register than people classified as less private.

Familiarity and Social Awareness

Respondent familiarity with the establishment, shown in Figure 3, did not have a significant impact on people's interest in playing music in the background (Q1), but it did reduce their interest in using headphones (Q2). Although familiarity was not a significant predictor of privacy (not shown), it did have a positive influence on how much people identified with their music (Q3). Furthermore, there was a link to the importance of background music in the establishment (Q5), which corresponds to the notion that music helps define a sense of place.

Familiarity had a strong correlation with how often people played music for other people in other situations (Q6), and how generally interested they were in relating to the music that was being played in the establishment (Q7, Q8). Furthermore, it had a highly significant impact on how interested people were in seeing information about *who* was playing music (Q9). On average, people reported that they were not interested in the people behind the music (Figure 1-A6), but factoring in familiarity, this activity becomes much more interesting to people (the only action significantly effected by this factor). Although they were interested in seeing information about others, familiar people were not significantly more likely to want to share information about themselves (data not shown).

From a design perspective, the role of familiarity represents the return customer, and as such affords many possibilities to reconcile the tensions concerning the management of personal information and identity in shared spaces. For example, a system could require regulars to register just once (thereby avoiding future hassles), and giving them, but not the general public, access to information about who is playing music. This design point represents an intermediate between an anonymous system where people can interact with the system without giving up any personal information, and a truly public system, where everybody can see the information relating to people who are using the system.

Ecosystem Stakeholders

Based on semi-structured interviews, *in situ* observations, and a literature survey, we have identified four major stakeholders that are necessary in order to support personal music in shared spaces: providers, contributors, proprietors, and listeners. Individual pieces of related work on music systems and the phone platform ([2], [8], [10], [14], [15], [17], [19], [23])each cover subsets of the overall ecosystem. When combined with perspectives gained from the interviews, a coherent view involving all the relevant positions emerges, leading to the concept of the shared space music lifecycle, diagramed in Figure 4.

The observations and interviews were conducted on/with patrons, owners, and employees at a neighborhood café and a local bar in the city of (Anonymized), chosen because of their progressive atmosphere, eclectic nature, and significant use of music. The café owner prides himself on running a "neighborhood" business, with a "regular" clientele. The music played is an eclectic mix of many different genres, and employees take turns selecting music to play from a collection of about 500 CDs provided by the owner. Employees are free to pick anything to play from the collection, and also to bring in CDs of their own "within reason." The staff follows guidelines about what kind of music to play when: "certain music is better to be played at certain times throughout the day, as the type of clientele and their moods change – it's ok to play more upbeat music later in the day." The owners of the local bar recently installed a Rockola digital jukebox [13]. Before the Rockola, the bar housed an old broken Jukebox, two televisions, and two video games. "We are not a music place," says the bar owner; however, since the Rockola was installed, music now plays constantly when there are about five or more people in the bar. "Since we installed the Rockola, nobody puts money in the video games. This is a big money maker for us."

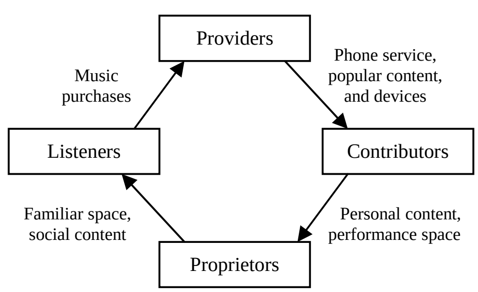

Fig. 4. The public music stakeholders and shared space music lifecycle. Each stakeholder contributes and receives different aspects of music while the circular link promotes commercial sustainability.

Stakeholder Organization

Each identified stakeholders has a specific role in the overall ecosystem:

- *Providers*: institutions that supply resources to catalyze the overall ecosystem. This category is really a collection of different organizations all provide different resources: music content, mobile music devices, or supporting cell-phone service.

- *Contributors*: individuals that personalize an environment, either by directly supplying digital music files for playback, or contributing their playlists or musical preferences to affect the music played in a space.
- *Proprietors*: institutions that consume resources (money from their patrons) and provide a shared context where contributors and listeners can come together to play/consume music. Proprietors are typically only indirectly involved in the music itself, often providing the space primarily for selling food or other consumables.
- *Listeners*: individuals that consume content, providing an audience. Like proprietors, they are often there for other reasons, such as for food or socializing, but they provide a critical component of the ecosystem.

Stakeholders manifest themselves slightly differently in different environments, shown in Table 1. In most traditional environments, some of the four roles are combined together into singular entities. This table is meant mainly as an illustration of how the various roles manifest themselves, and does not provide a concrete enumeration of the roles in the different contexts. Parallels for applying Musicology to a home environment, instead of public places, can easily be seen in the diagram (Table 1), where the home owner takes on the role of proprietor.

Empowering the contributor to affect music played in a space is the key concept behind Musicology: enabling individuals to contribute their personal content by removing the technology and interaction barriers presented by existing music systems. The providers hold a very important position in the public music ecosystem because they control the influx of resources, an aspect not highlighted by other public music frameworks [15][17]. Likewise, compared with general cell-phone ecosystems [19], the listener is provides the audience and cycle back to the providers, which is the music equivalent to the person on the other end of a telephone call.

Enabling the Stakeholder Lifecycle

Each stakeholder has various needs, preferences, and constraints that ultimately define the stakeholder lifecycle. Interaction within and between stakeholders is shaped by

<table>
<tr><th>Environment\Stakeholder:</th><th>Provider</th><th>Contributor</th><th>Proprietor</th><th>Listener</th></tr>
<tr><td>Mobile Music Player:</td><td>Various</td><td colspan="3">The Individual</td></tr>
<tr><td>Trad. Public Place:</td><td>Various</td><td colspan="2">Store Staff</td><td>Patrons</td></tr>
<tr><td>House Party:</td><td>Various</td><td colspan="2">Home Owner</td><td>Guests</td></tr>
<tr><td>Live Concert:</td><td colspan="3">Performer</td><td>Audience</td></tr>
<tr><td>Music Festival:</td><td colspan="2">Stage Performers</td><td>Festival Operator</td><td>Attendees</td></tr>
<tr><td>Musicology in Café:</td><td>Various</td><td>Individual</td><td>Store Staff</td><td>Individual</td></tr>
<tr><td>Musicology at Home:</td><td>Various</td><td>Guests</td><td>Home Owner</td><td>Guests</td></tr>
</table>

Table 1. Stakeholders in a variety of different contexts. The Provider stakeholder is listed as “various” for simplicity since it always encompasses a variety of different providers (content, service, equipment, device, etc…).

their defining feature, and understanding these features is necessary to ensure that the needs of all stakeholders are met when designing a music system for shared spaces.

*Provider*s are primarily concerned with **compensation** for their services and product **differentiation**. From a pragmatic perspective, they need to realize a return on their investment in order to supply future services. A mobile device provider, for example, will want to receive revenue for selling their device and be able to differentiate their device from others on the market. The complexity behind the providers is exemplified by the release of the Rokr mobile music phone [10], which was a joint announcement between a wireless carrier (Cingular), mobile phone manufacturer (Motorola), and music source (Apple).

Contributors primarily have a need for **identity management**, as discussed in the previous section. This underlying need drives their acquisition of resources from the providers, causing them to acquire new content, seek out newer phones, etc… Likewise, they will also seek out appropriate venues in which to express their musical tastes (such as a technology savvy café), creating a link to proprietors. Contributors share a dual role with listeners, depending on when/if *their* song is playing at the time.

Proprietors are mainly interested in **cultivating** a desirable environment and in **protecting** themselves from litigation regarding the improper use of copyrighted content. Cultivating the environment requires balancing the perspectives of the contributors, proprietors, and listeners to effectively control what music is played music in the space. The value proposition for proprietors is demonstrated by the rapid adoption of free WiFi hot-spots at many cafés in order to attract customers, despite the fact that customers using laptop computers often sit there for a long time.

Listeners are primarily focused on **enjoyment** and possibly **capture** elements of their experiences**.** For example, if they hear a song they particularly enjoy, they can to capture the experience by recording the name of the artist and the title of the song. Ultimately, the providers can benefit from this experience capture, as the listeners will potentially turn to the provider to acquire new content.

Implicit in some of the concerns associated with different stakeholders are some difficult challenges, such as copyright protection, deploying technologies in the face of chicken-and-egg dynamics, and a balance between contributor and proprietor control over the music played in a space. The implication of these barriers is discussed further in the Adopting the Ecosystem section, below.

Supporting Technology

A comparison between two systems supporting wireless music playback, one using a basic implementation and another using more capable protocols, highlights how existing technologies can be adapted to better support social music interaction. Bluetooth, used by both systems, the dominant wireless technology found in handheld devices and can easily support the data

Fig. 5. The Motorola e680i multimedia cell-phone platform, supporting both A2DP- and PAN- based wireless streaming over Bluetooth.

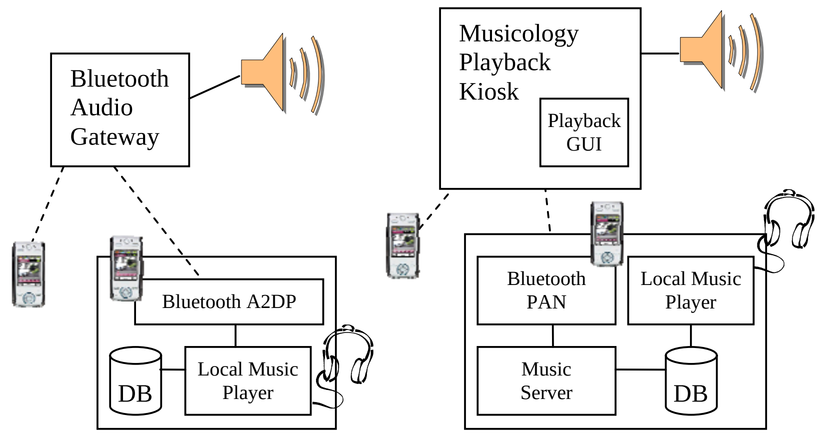

Fig. 6. Two different architectures supporting wireless music playback. On the left, an out-of-the-box implementation based on the Bluetooth A2DP profile. On the right, a prototype solution based on the Bluetooth PAN profile. The music database (DB) represents the user's local music store, and the Local Music Player is the native application designed to play music files locally (e.g., through headphones).

rates and ranges necessary to support wireless music. Many advanced smart-phone devices already posses enough storage and processing capability to act as mobile music players. This comparison uses the Motorola e680i, a high-end "multimedia phone" platform [12], shown in Figure 5, based on the Intel® XScale™ processor with built-in Bluetooth and a SD/MM card slot for removable storage. The device runs the Linux operating system and has a built-in media player capable of playing mp3 files, as well as video. The advanced processing capability, general-purpose OS, and removable storage card enabling the platform to support complex applications which have traditionally only been associated only with laptop or desktop systems: Specifically, the system is capable enough to run DLNA-compliant media servers over a TCP/IP network.

This investigation focuses specifically on the system components necessary to enable multiple devices effectively connecting and providing content for a shared-music ecosystem. Inherent in the Bluetooth protocol is a discovery and connection mechanism that allows devices to discover each other and form peer-to-peer connections: for example, from a mobile phone one can discover and connect to a wireless headset for listening to music, or one could discover and connect to music access point in the environment. This connection process is fundamentally the same for both implementations, since they both use Bluetooth. The iRadio system, discussed previously in related work, is based on the A2DP implementation, thereby inheriting all of its advantages and disadvantages.

A2DP and PAN Implementations

The two implementations for wireless playback on the e680i device differ in the basic standard upon which they are built: Advanced Audio Distribution Profile

(A2DP) vs. Personal Area Networking (PAN). Additionally, a local playback mechanism, which plays music through wired headphones, is used as a baseline comparison. The A2DP design is based on the basic cable replacement model, where the wireless link is used to stream content directly to a wireless stereo headset, conceptually replacing the cable between handset and headset. The PAN implementation is based on a more general network model, where devices join the network, like they would over wired Ethernet, and use TCP/IP based communication. The advantage of the A2DP model is simplicity, while the PAN model offers flexibility. The resulting system architectures, diagramed in Figure 6, reflect this distinction, where A2DP is easier to implement, but PAN supports more flexible access and control.

The A2DP profile specifies protocols and formats for transferring music across a wireless link. Primarily intended for devices such as wireless headsets, this profile provides a set of "lowest common dominator" audio formats designed to be easy to implement (reducing the cost of the headsets); however, in order to send audio data, which is typically stored in MP3 or some other highly compressed format, the platform must transcoded the data on the fly, increasing the energy consumption of the device. On the "remote" side, the A2DP system connects to any compatible audio gateway. The underlying A2DP profile itself is capable of supporting arbitrary format encodings, but it requires that both endpoints support the capability, and MP3 or other highly-compressed formats are not part of the standard. A2DP does not inherently support multiple connections over the same stream, but basic Bluetooth multiple-device support can be used on the audio gateway to set up multiple independent A2DP streams between multiple devices, and devices can simply be paused when not active.

The Bluetooth Audio/Video Remote Control Profile (AVRCP) is used as the control mechanism across an A2DP channel. AVRCP allows basic remote control operations that one would find on a simple audio interface: e.g., play, pause, next song, and volume control. Neither A2DP nor AVRCP allows arbitrary access to a user's music library – instead, only basic navigation within an already playing stream is allowed. So, in order to stream music wirelessly from a mobile device, the device itself must first start playing the song locally, and then A2DP can be used to stream the audio off-device, and AVRCP can be used to control it remotely.

To fully realize the capabilities of the underlying platform, we built the PAN-based prototype system implementing a custom music server using standard TCP/IP sockets. Audio is communicated using raw transfers of the MP3 (or other format) files, pushing all decompression activities to the playback system. This model is different than the A2DP implementation above, which transcodes files locally to a simpler format for communication: the PAN model is less resource intensive on the mobile device, but *more* resource intensive on the playing end – the ramifications of which are discussed further in the next section. On the receiving end, the PAN system uses a custom Java-based music player running on a general purpose Windows XP machine. Functionally similar to any standard desktop music software, such as iTunes of Music match, the custom implementation allows underlying Bluetooth connection mechanics to be hidden from the user. In order to support multiple connections, the PAN model can simply allow multiple devices to connect at the same time and contact them using their respective IP addresses.

Ironically, the PAN implementation requires more *functionally* to be supported on the mobile device, in terms of network stack and server capabilities, but is less *resource* intensive because incurs lower CPU utilization. The A2DP implementation is architecturally easier to implement, but requires more local CPU utilization to operate. The fundamental requirements here are driven by the playback side of the system, not the mobile device: A2DP is designed to work with low-capability headsets, while PAN is designed to work well on high-capability systems. Furthermore, from an implementation perspective, the A2DP solution is easier to integrate with existing software solutions because it can transparently displace the operating system's audio driver, while PAN requires a customized applications to communicate over the network socket. The difference between the two architectures can be seen in Figure 6 by the location of the Bluetooth stack in relation to the local music player.

In summary, there are two fundamental differences between the architectures described above that have a significant impact on remote wireless music playback:

1. A2DP transcodes and streams audio across the link, minimizing the resource requirements on the receiving end, while PAN transfers files en-masse, minimizing the resource utilization on the sending end.
2. PAN allows arbitrary access to a device's music library, while A2DP can only get content from the current stream, affecting playback with simple remote control commands such as start/stop/forward/back.

Power and Control

Two aspects of the implementations above directly effect wireless playback: power and control. In short, the "simpler" A2DP implementation consumed more power on the mobile device because it must do more local processing. Furthermore, the PAN model affords more flexible user interfaces since it is built on more common network standards.

The energy consumption profile for the various implementations, Figure 7, shows the relative energy efficiency of the various systems. Local playback represents the power consumption of the device playing the audio locally through headphones. The A2DP Streaming test streams data between the e680i and a Motorola DC800 home audio gateway, while the PAN Transfer test transfers the music file between the e680i and an IBM ThinkPad laptop. Power measurement is accomplished using a technique similar to [21], where the voltage drop across a sense resistor placed between the device and battery is measured to calculate

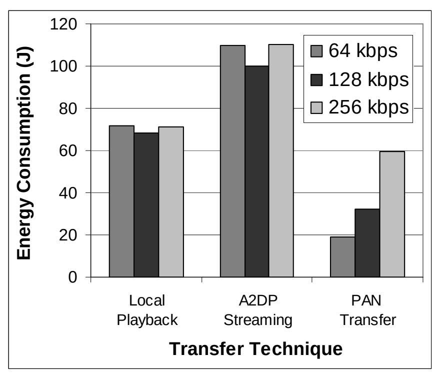

Fig. 7. The system energy consumption for a variety of playback techniques. The music source is a 172s MP3 file encoded at a variety of bitrates. The energy measurements encompass the entire platform, with the backlight turned off. The PAN transfer technique takes 35s, 67s, and 128s to complete, depending on the underlying bitrate.

power. Combining the power consumption with the time required to complete the operation, which is the duration of the song for Local and A2DP, but just the time required to transfer the file for PAN, yields the energy required to play the given song – measuring the battery lifetime impact on the mobile device.

The energy difference between Local Playback and A2DP Streaming represents the difference between driving the local power amplifier to power the headphones, and transferring the audio across the wireless link; both techniques use a local application to convert the native MP3 files into another format – an inherently CPU intensive task. Since neither the Local Playback nor A2DP Streaming techniques can process the entire music file at one time, the system must remain active for the duration of the song in order to maintain playback. The PAN technique, on the other hand, simply transfers the source file off the device as quick as it can, allowing the device to return to a low-power sleep mode; the recipient of the file can start playing it as soon as the first packet arrives, maintaining a quick playback response. The average power consumption of the transfer technique (490 mW) is higher than during local playback (409 mW), but the reduced activation time results in lower overall energy consumption. The A2DP solution realizes the worst of both worlds: requiring the system to remain active for the duration of the song as well as having higher average power consumption (620 mW).

In essence, the PAN implementation is turning the mobile device into a general purpose computing platform – equivalent in many ways to a laptop or desktop computer. Although the music server implementation on the phone is a custom-written application, it could just as easily be a port of iTunes or some other standard music library and server manager system, demonstrating that the mobile platform, using the appropriate standards, has the capability to be integrated in already established media sharing systems (such as the iTunes sharing network described under Related Work).

Since the PAN software implementation runs on the same hardware as the A2DP technique, the additional cost to implement a truly flexible wireless music platform is minimal. Furthermore, the power and control deficiencies inherent with the A2DP model are not conducive to a full-featured wireless music solution. One concern with the more flexible PAN model is security: a more open networking standard is potentially more vulnerable to virus attacks. Fortunately, the underlying Linux OS used in the e680i device is virtually the same as found in common desktop and server systems, allowing it to be hardened against viruses and other security threats using standard techniques such as firewalls, although this is an area that still requires future work to test and validate the security proposition of these mobile devices.

Evolution and Adoption of Music Technologies for Shared Spaces

The real-world adoption of a Musicology-like system faces a chicken-and-egg problem since it relies on both installed infrastructure and personal mobile devices. What impetus do providers have to create devices and services that will supplant existing practices for playing music in shared spaces? What will encourage proprietors to install infrastructure if there are no users out there to use it? Other music technologies, like stand-alone mp3 players, online music store, standalone jukebox systems, etc… are point solutions that don't require extensive integration across multiple domains. Even the relatively simple integration of basic music player capabilities into a cell-phone device is a very complex proposition because it involves multiple providers that are all trying to benefit commercially from the system.

One simple way to encourage a subset of the Musicology vision is to incrementally enhance laptop-based music sharing models, circumventing the need for required modifications to cell-phone platforms, which are inherently closed by the manufacturers. The freely available MusicMatch music player software, for example, can be configured to be a DLNA-compatible music server; then, since it is compliant with the open DLNA standards, one could write a digital jukebox application that would scan the network and automatically play music from any available server (as long as the source files are stored in a non-copy-protected format). This is similar to the current iTunes sharing models, but would seamlessly integrate multiple people into a single space. Alternatively, one could set up a publicly-available AirTunes audio gateway, allowing anybody attached to a local WiFi-hotpot to push their music into the shared space, using a socially mediated shared access protocol where people manually disconnect their computer since AirTunes only allows one device to play music at a time. Although not involving cell-phone class devices, these implementations would create a public shared space for music, exploring the social viability of such a system.

Enabling mobile devices for music playback in a private home environment is likely to occur well before public environments, primarily because one person can act as both a contributor and a proprietor and enable their home to play music from their devices. The basic capabilities to enable communication with DLNA-compatible mobile phones are currently being incorporated into future digital home standards, and in a few years, compatible phones will start to reach the market. This capability, combined with a multi-user shared playback system, similar to one described in the preceding paragraph, will enable mobile phones to provide basic shared space playback capability. This solution is not likely to afford the same level of control as a fully integrated system, but it will jumpstart the basic capability and draw the attention of providers to work on better integrating their solutions.

The concern over enforcement of DRM in public spaces is a significant barrier for adoption since many proprietors will hesitate to adopt public playback systems for fear of litigation. Fortunately, the needs of DRM solutions for public spaces are partially addressed by a combination of necessary solutions for private systems and public licenses that are required for the public playback of recorded music. Basically, it is not strictly legal to play recorded music publicly (e.g., a café owner playing personal CDs in public), without acquiring a performance license from the appropriate copyright holder. Any café environment that plays "personal" CDs should have such a license, and the café investigated in the ecosystem section in fact pays two separate fees to cover their CD collection. However, there is still an underlying matching problem: It is necessary to ensure that a host system will only play music for which it has an appropriate public license, not just music for which there is a private license. One solution to this problem would be to use an internet-accessible database to look up the rights holder for the song (unless the info is already included in song's metadata), and then ensure that the establishment holds the proper performance rights.

A big issue facing the adoption of personal music into shared public places is that of control by the proprietor: how can they allow individuals to contribute music while still maintaining enough control to make the listening environment acceptable to the majority of their customers. This was a concern highlighted in the informational interviews described earlier, as well as the Jukola project. Ideally, the management could allow a fairly wide range of music choices while preventing anything that will not be appreciated by a majority of the listeners. The solution offered by traditional jukebox systems, and the Jukola model, is to provide a relatively limited selection of music for patrons to choose from – a solution that does not work if truly personal music is to be allowed in the space. In order to vet the appropriateness of arbitrary personal music content, an internet-accessible database providing sufficient information to allow a proprietor to decide which songs can be played could be used: a contributor would submit a song for playback, and the host system would look up the information about the song and play it if the song if it met the pre-set criteria.

None of these aspects address the user-interface design on the mobile device, which is an important component of the music cycle (Figure 4). The UI design enables both the contributors and the listeners to seamlessly interact with the system, while also providing a path for listeners to easily discover new music. Fortunately, a new user interface is not strictly necessary to build and deploy a rudimentary system, based on the possibilities described in the preceding paragraphs. Then, once rudimentary playback systems are in place, there will be a pull to develop applications for mobile platforms that support a better user experience, thus closing the cycle.

Conclusion

The Musicology system has provided valuable insights into the issues influencing personal music playback in shared public spaces. Musicology is primarily based on using wireless personal devices such as cell-phones to bring music content into cafés, the Digital Home, and other social environments. Wireless laptop computers can also serve as sources of personal content for this use. User survey's investigating this concept highlight the role of privacy as an identity boundary, as well as the impact that patron familiarity has on a space. Furthermore, an analysis of the supporting ecosystem details the complex relationships between players, and sheds some light on how DRM technologies and control issues may affect the use of music in public places. Mobile handsets and the necessary supporting infrastructure is starting to exists, but integrated solutions are still necessary to enable the overall vision.

By studying the concrete we can better understand the abstract. Music is the first step in a long line of media, and ultimately actions, that can be sourced from a user's personal device. The concepts developed here can easily be applied to photographs and video, enabling a new channel for people to express themselves and customize physical spaces – such expression is already rampant in the form of blogs, photo sharing web-sites, and repositories of personal videos, and our personal mobile devices have the capability to brings them directly to the real world. As the supporting infrastructure evolves, personal customization will evolve with it, allowing us to be better connected with our local environment and those around us, instead of limiting our interaction through isolated islands of computation.

References

[1] Barkhuus, L., Dey, A.: Location-Based Services for Mobile Telephony: a study of users' privacy concerns. In: Proc. Interact 2003 (2003) 709 -71.

[2] Bassoli A, Moore J., and Agamanolis S., tunA: Socialising Music Sharing on the Move (book chapter), in Kenton O'Hara and Barry Brown (eds), Consuming Music Together: Social and Collaborative Aspects of Music Consumption Technologies

[3] Brown, B., A. Sellen, E. Geelhoed (2001) Music sharing as a computer supported collaborative application. In: Proceedings of ECSCW 2001, Bonn, Germany. Brown, B., Geelhoed, E. & Sellen, A. (2001), The Use of Conventional and New Music Media: Implications for Future Technologies, Proc. of INTERACT'01, Tokyo, Japan.

[4] Bull, M.: Investigating The Culture of Mobile Listening: From Walkman to IPod, from Consuming Music Together, Book chapter), ISBN 1402040318, Feb 28, 2006

[5] De Nora, T. (1986) How is extra-musical meaning possible. Music as a place and space for work. Sociological theory 4, 84-94.

[6] De Nora, T. (2000) Music in Everyday Life. Cambridge: Cambridge University Press

[7] Firth, S. and A. Goodwin, 1990, "From subcultural to cultural studies", in S. Firth and A. Goodwin (Eds.) On record: Pantheon Books.

[8] http://broadband.motorola.com/iradio/

[9] http://www.apple.com/airportexpress/

[10] http://www.apple.com/itunes/mobile/

[11] http://www.dlna.org/home

[12] http://www.motorola.com/motoinfo/product/details/0,,103,00.html

[13] http://www.rock-ola.com/

[14] Jackson, M. , Singh, S. , Beekhuyzen J., Waycott J., DRMs, fair use and users' experience of sharing music, Proceedings of the 5th ACM workshop on Digital rights management, November 07-07, 2005, Alexandria, VA, USA

[15] McCarthy, J. F., and. Anagnost. T. D. (1998). MUSICFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts. In Proceedings of the ACM 1998Conference on Computer Supported Cooperative Work (CSCW '98), Seattle

[16] Mobile Music: Ringtones, Ringbacks & Full Tracks; Juniper Research Limited, February 1, 2005 – Pub ID: JUN1083825

[17] O'Hara, K., Lipson, M., Jansen, M., Unger, A., Jeffries, H., Macer, P.: Jukola: democratic music choice in a public space. Conference on Designing Interactive Systems 2004

[18] Palen, L. and Dourish P. Unpacking "privacy" for a networked world. In Proc. CHI 2003, CHI Letters 5(1), 129-136.

[19] Palen, Leysia, Salzman, Marilyn C. (2002): Beyond the handset: designing for wireless communications usability. In ACM Transactions on Computer-Human Interaction

[20] Privacy & American Business: Consumer Privacy Attitudes: A Major Shift Since 2000 and Why. 10 (2003)

[21] Raghunathan, V., Pering, T., Want, R., Nguyen, A., Jensen, P.: Experience with a low power wireless mobile computing platform. ISLPED 2004: 363-368

[22] Smith, I., Consolvo S., LaMarca A., Hightower, J., Scott J., Sohn, S., Hughes J., Iachello, G., Abowd, G.: Social Disclosure of Place: From Location Technology to Communication Practices. Pervasive 2005: 134-151

[23] Voida, A.; Grinter, R. E.; Ducheneaut, N.; Edwards, W. K.; Newman, M. Listening in: practices surrounding iTunes music sharing. Proceedings of ACM Conference on Human Factors in Computing Systems (CHI 2005); 2005 April 2-7; Portland; OR; USA. NY

[24] Want R., Pering T., Danneels G., Kumar K., Sundar M., and Light J.. The Personal Server: Changing the Way We Think about Ubiquitous Computing. In Proceedings of UbiComp 2002.

[25] Weiser, M., The computer for the 21st century, Scientific American (September 1991)