

# Large-Scale Mobile Audio Environments for Collaborative Musical Interaction

Mike Wozniowski &  
Nicolas Bouillot  
McGill University  
Centre for Intelligent Machines  
Montréal, Québec, Canada  
{mikewoz,nicolas}@cim.mcgill.ca

Zack Settel  
Université de Montréal  
Montréal, Québec, Canada  
zs@sympatico.ca

Jeremy R. Cooperstock  
McGill University  
Centre for Intelligent Machines  
Montréal, Québec, Canada  
jer@cim.mcgill.ca

## ABSTRACT

By freeing users from the confines of computer terminals and interfaces that severely limit mobility, new application spaces and artistic forms emerge, where interaction is driven by user movement. We present a system to manage large-scale mobile audio environments, where multiple participants share an overlaid continuous virtual audio environment, while navigating through real physical space. Each user is equipped with a mobile computing device, GPS receiver, orientation sensor, microphone, headphones, or various combinations of these technologies. We investigate methods of location tracking, wireless audio streaming, and state management between mobile devices and centralized servers. The result is a system that allows mobile users, with subjective 3-D audio rendering, to share virtual scenes containing interactive audio elements. These elements can be organized into large-scale spatial audio interfaces, thus allowing for immersive mobile performance, locative audio installations, and many new forms of collaborative sonic activity.

## Keywords

sonic navigation, mobile music, spatial interaction, wireless audio streaming, locative media, collaborative interfaces

## 1. INTRODUCTION & BACKGROUND

With the design of new interfaces for musical expression, it is often argued that control paradigms should capitalize on natural human skills and activities. As a result, a wide range of tracking solutions and sensing platforms have been explored, which translate human action into signals that can be used for the control of music and other forms of media. The physical organization of interface components plays an important role in the usability of the system, since user motion naturally provides kinesthetic feedback, allowing a user to better remember the style of interaction and gestures required to trigger certain events. As digital devices become

increasingly mobile and ubiquitous, we also expect interactive applications to become more distributed and integrated within our physical environment. This prospect yields a new domain for musical interaction employing augmented-reality interfaces and large multi-user environments.

We present a system where multiple participants can navigate about a university campus, several city blocks, or an even larger space. Equipped with position-tracking and orientation-sensing technology, their locations are relayed to other participants and to any servers that are managing the current state. With a mobile device for communication, users are able to interact with an overlaid virtual audio environment, containing a number of processing elements. The physical space thus becomes a collaborative augmented-reality environment where immersive musical interfaces can be explored. Musicians can input audio at their locations, while virtual effects processors can be scattered through the scene to transform those signals. All users, performers and audience alike, receive subjectively rendered spatial audio corresponding to their particular locations, allowing for unique experiences that are not possible in traditional music performance venues.



Figure 1: A mobile performer

### 1.1 The Audioscape Engine

In order to facilitate the creation and management of augmented-reality scenes, we use our *Audioscape* engine [1], which allows for the spatial organization of sound processing, and provides an audiovisual rendering of the scene for feedback. Audio elements can be arranged in a 3-D world and precise control over the directivity of propagating audio

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

NIME08, Genova, Italy

Copyright 2008 Copyright remains with the author(s).

is provided to the user. For example, an audio signal emitted by a sound generator may be *steered* toward an effects processor that exists at some 3-D location [25]. The processed signal may again be steered towards a virtual microphone that captures and busses the sound to a loudspeaker where it is heard. The result is a technique of spatial signal bussing, which lends particularly well to many common mixing operations. Gain control, for instance, is accomplished by adjusting the distance between two nodes in the scene, while filter parameters can be controlled by changing orientations.

A large portion of our work has been the investigation of methods for controlling the position of a user within the virtual scene. Early versions of our system operated mainly with game-like techniques, where users stood in front of screens, and navigated through a scene using controllers such as joysticks or gamepads. Audio was spatially rendered according to the user's current virtual position and orientation (viewpoint), using multiple loudspeakers or headphones. The first step towards physically involving users with the virtual 3-D audio content around them was to couple a head-mounted orientation sensor with a headphone audio display. Thus, a user's real head movement would yield corresponding changes to the apparent spatial audio rendering about them. To further extend this degree of physical involvement we began to add real-world location/position awareness to the system, allowing users to move around the space physically instead of using input devices such as joysticks or mice to indicate their position.

We have already experimented with small-scale motion tracking, using computer vision algorithms from the OpenCV library [5]. For example, the *4Dmix*<sup>3</sup> installation [7] tracked up to six users in an 80m<sup>2</sup> gallery space. The motion of each user controlled the position of a recording buffer, which could travel among a number of virtual sound generators in the scene. The result was a type of remixing application, where users could control the mix by moving through space. Taking this one step further, we now include larger-scale position tracking, such as that of a Global Positioning System (GPS). The great appeal of wide-area locative technologies is in the solutions they bring to multi-user immersive audio and music applications, and in the implicit conceptual leap of scale that translates directly to new musical experiences, and thus new forms.

## 1.2 Mobile Music Venues

Freeing users from the confines of computer terminals and interfaces that severely limit mobility, application spaces emerge that can operate in a potentially unbounded physical space. These spaces offer many novel possibilities that can lead to new artistic approaches; or they can re-contextualize existing concepts that can then be revisited and expanded upon. An excellent example is parade music, where sound emission is spatially dynamic or mobile; a passive listener remains in one place while different music is coming *and* going. One hundred years ago, Charles Ives integrated this concept into symphonic works, where different musical material *flowed* through the score, extending our notions of counterpoint to include those based on musical material *proximity*. The example of parade music listening expands to include two other cases: a mobile listener can walk with or against the parade, yielding additional musical experiences. Our work also integrates the concept of active listening; musical material may be organized topographically

in space, produced by mobile performers and encountered nonlinearly by mobile listeners. From this approach come several kinds of rich musical forms, which like sculpture, integrate *point of view*; listeners/observers create their own unique rendering of the musical experience. Thus, artists may create works that explore the spatial dynamics of musical experience, where *flowing* musical material is put in counterpoint by navigation. Musical scores begin to resemble maps, and listeners play a larger role in authoring their experiences.

## 1.3 Related Work

Some projects have considered the task of tagging geographical locations with sound. The [*murmur*] project [4] is one simple example, where users tag interesting locations with phone numbers. Others can call the numbers using their mobile phones and listen to audio recordings related to the locations. Similarly, the *Hear&There* project [19] provides users with the ability to record audio at a given GPS coordinate, while providing a spatial rendering of other recordings as they walk around. Unfortunately, this is limited to a single-person experience, where the state of the augmented reality scene is only available on one computer. Tanaka proposed an ad-hoc (peer-to-peer) wireless networking strategy to allow multiple musicians to simultaneously share sound using hand held computers [21]. Although the system does not incorporate position awareness, other work by Tanaka and Gemeinboeck [22] capitalized on location-based services (LBS) available on 3G cellular networks to acquire coarse locations of mobile devices. They propose the creation of *locative media instruments*, where geographic localization is used as a musical interface.

Large areas can also be used for musical interaction in other ways. *Sonic City* [15], for example, proposes mobility for interaction rather than just location. As a user walks around a city, urban sounds are processed in real time as a result of readings from devices such as accelerometers, light sensors, temperature sensors, and metal detectors. Orientation or heading information in particular can provide useful feedback, since spatial sound conveys a great deal of information about directions of objects and the acoustics of an environment. Projects including *GpsTunes* [20] and *Melodious Walkabout* [14] use this type of information to provide audio cues that guide individuals in specific directions.

We take inspiration from the the projects mentioned above, and incorporate many of these ideas into our work. However, real-time high-fidelity audio support for multiple individuals has not been well addressed. Tanaka's work [21], as well as some of our past experiences [9], demonstrate how we can deal with the latencies associated with distributed audio performance, but minimizing latency remains a major focus of our work. The ability to create virtual audio scenes is easily supported by our existing Audioscape engine. To address the further needs of distributed mobile interaction, we are adding large-scale location sensing and the ability to distribute state, signals, and computation among mobile clients effectively. These challenges are addressed in the following sections.

## 2. LOCATIVE TECHNOLOGY

In order to support interaction in large-scale spaces, we require methods of tracking users and communicating between them. A variety of mobile devices are available for this pur-

pose, potentially equipped with powerful processors, wireless transmission, and various sensing technologies. For our initial prototypes, we chose to develop on Gumstix (verdex XM4-bt) processors with expansion boards for audio I/O, GPS, storage, and WiFi communication [17]. These devices have the benefit of being full-function miniature computers (FFMC) with a large development community, and as a result most libraries and drivers can be supported easily.

## 2.1 Wireless Standards

In other recent work [10], we offer a comparison of different wireless technologies. We note, however, that Bluetooth and WiFi standards are the most generally available on mobile devices. For transmission of data between sensors located on the body and the main processing device, Bluetooth is a viable solution. However, even with Bluetooth 2.0, a practical transfer rate is typically limited to approximately 2.1 Mbps. If we want to send or receive audio (16 bit samples at 44kHz), approximately 700 kbps of bandwidth is needed for each stream. In theory, this allows for interaction between up to three individuals, where each user sends one stream and receives two. Given the need to support a greater number of participants, we are forced to use WiFi.<sup>1</sup> Furthermore, the range of Bluetooth is limiting, whereas WiFi can relay signals through access points. Furthermore, we can make use of higher-level protocols such as Optimized Link State Routing protocol (OLSR) [18], which computes optimal forwarding paths for ad-hoc nodes. This is a viable way to reconfigure wireless networks if individuals are moving.

## 2.2 GPS

GPS has seen widespread integration into a variety of commodity hardware such as cell phones and PDAs. The position of these devices can thus be tracked in outdoor environments, and provide the 3-D geospatial coordinates of users. However, GPS accuracy in consumer-grade devices is quite poor, starting at about 5m in the best case (high-quality receiver with open skies) [24], and ranging possibly up to 100 metres [8]. There are, however, some methods that reduce this error. For example, differential GPS (DGPS) uses carefully calibrated base stations that transmit error corrections over radio frequencies. The idea is that mobile GPS units in the area will have similar positional drift, and correcting this can yield accuracies of 1m. Another technique, known as assisted GPS (AGPS), takes advantage of common wireless networks (cellular, bluetooth, WiFi) in urban environments to access reference stations with a clear view of the sky (e.g., on the roofs of buildings). Although accuracy is still in the order of 15m, the interesting benefit of this system is that localization can be attained indoors (with an accuracy of approximately 50m) [8].

## 2.3 Orientation & Localization

While GPS devices provide location information, it is also important to capture a listener's head orientation so that spatial cues can be provided, the resulting sound appearing to propagate from a particular direction. Most automotive GPS receivers report heading information by tracking the vehicle trajectory over time. This is a viable strategy for in-

ferring the orientation of a vehicle, but a listener's head can change orientation independently of body motion. Moreover, the types of applications we are targeting will likely involve periods of time where a user does not change position, but stays in one place and orients his or her body in various directions. Therefore, additional orientation sensing seems to be a requirement.

In human psychoacoustic perception, the accuracy and responsiveness of orientation information is important, since a listener's ability to localize sound is highly dependant on changes in phase, amplitude, and spectral content with respect to head motion. Responsiveness, in particular, is a significant challenge, considering the wireless nature of the system. Listeners will be continuously moving their heads to help localize sounds, and a delay of more than 70ms in spatial cues can hinder sound localization ability [11]. Furthermore, it has been demonstrated that head tracker latency is most noticeable in augmented reality applications, since a listener can compare virtual sounds to reference sounds in the real environment. In these cases, latencies as low as 25ms can be detected, and begin to impair performance in localization tasks at slightly greater values [12]. For this reason, it is suggested that latency is maintained below 30ms.

To track head orientation, we attach an inertial measurement unit (IMU) to the headphones of each participant, capable of sensing instantaneous 3-D orientation with an error of less than 1 degree. It should be mentioned that not all applications will require this degree of precision, and some deployments could potentially make use of trajectory-based orientation information. For instance, the Melodious Walkabout [14] uses aggregated GPS data to determine the direction of travel, and provides auditory cues to guide individuals in specific directions. Users hear music to their left if they are meant to take a left turn, whereas a low-pass filtered version of their audio is heard if they are travelling in the wrong direction. We can conceive of other types of applications, where instantaneous head orientation is not needed, and users could adjust to the paradigm of hearing audio spatialization according to trajectory rather than line of sight. Of particular interest, are high-velocity applications such as skiing or cycling, where users are generally looking forward, in the direction of travel. Such constraints can help with predictions of possible orientations, while the faster speed helps to overcome the coarse resolution of current GPS technology.

## 3. WIRELESS AUDIO STREAMING

The move to mobile technology presents significant design challenges in the domain of audio transmission, largely related to scalability and the effects of latency on user experience. The use of WiFi networks means less bandwidth and more instability than traditional wired networks, and the musical context of our application requires a real-time synchronization that is difficult to achieve with present transmission methods. Furthermore, a certain level of quality needs to be maintained to ensure that mobile performers and audience members experience audio fidelity that is comparable to traditional venues. Satisfying these requirements while maintaining the emphasis on cheap, small and lightweight devices as demanded by the market remains a notable challenge.

### 3.1 Latency

<sup>1</sup>We note viable alternatives on the horizon, such as the newly announced SOUNDabout Lossless codec, which allows even smaller audio streams to be sent over Bluetooth.

Mobile applications tend to rely on compression algorithms to respect bandwidth constraints, and as a result they often incur signal delays that challenge musical interaction and performer synchronization. Acceptable latency tolerance depends on the style of music, with figures as low as 10ms [13] for fast pieces. More typically, musicians have difficulty synchronizing with latencies above 50ms [?]. Most audio codecs require greater than this amount of encoding time.<sup>2</sup> Due in part to limited computational resources available on our mobile devices, we instead transmit uncompressed audio, thus fully avoiding codec delays in the system.

Other sources of latency include packetization delay, corresponding to the time required to fill a packet with data samples for transmission, and network delay, which varies according to network load and results in jitter at the receiver. Soundcard latencies also play a role, but we consider this to be outside of our control. The most effective method for managing these remaining delays may be to minimize the size of transmitted packets. By sending a smaller number of audio samples in each network packet, we also decrease the amount of time that we must wait for those samples to arrive from the soundcard. However, there is a tradeoff between decreasing packet size and the resulting increase in the number of packets on the network. As the packet rate increases, the bandwidth also increased due to the extra packet headers, and the network can become unstable resulting in high jitter.

On the receiving end, audio samples are typically held in a playback buffer, designed to maintain a constant latency and thus mask the effects of late arriving packets. While we can decrease the size of this receiving buffer, care must be taken to maintain quality in the case of high jitter situations.

### 3.2 Scalability

Eliminating the use of compression risks an explosion of bandwidth as the number of users grows, since there is potential for an  $n^2$  effect of audio interactions between them. We have investigated a number of solutions to deal with potential bottlenecks. Multicast technology, for instance, allows devices to send UDP packets to an IP multicast address. Interested clients are able to subscribe to the streams of relevance, drastically reducing the overall required bandwidth. However, our tests with several 802.11 network interfaces revealed very poor multicast support through their Linux drivers. Many were unable to support even a modest throughput of 700 kbps, exhibiting data loss of 50% at this level.

In a related publication [10], we collected statistics for various chipsets and found that multicast transmission experienced higher jitter than unicast, mandating a larger receiver buffer to maintain quality. Furthermore, the packet loss for the multicast tests was in the order of 10-15%, resulting in a distorted audio stream while unicast had almost negligible losses of 0.3%. From these findings, we have decided to await improved multicast support in wireless chipsets, and rely for now on a point-to-point streaming methodology.

### 3.3 Adaptive Streaming

To deal with the bandwidth limitations of WiFi networks and the lack of reliable multicast technology, we have for-

<sup>2</sup>Possible exceptions are the Fraunhofer Ultra-Low Delay Codec (offering a 6ms algorithmic delay) [23] and the SOUNDabout Lossless codec (claiming under 10ms).

ulated a method to adapt the content of streamed audio in order to maximize the number of users that can be supported by the system. Accordingly, we have developed an adaptive transmission protocol [10] for high fidelity audio, called *nstream*, which supports dynamic adjustment of sender throughput and receiver buffer management. This is accomplished by switching between different levels of PCM quantization (8, 16 and 32 bit), packet size, and receiver buffer size. The protocol is developed as an external for Pure Data [6], and can be deployed on both a central server and a mobile device.

In benchmark tests, we have successfully transmitted uncompressed streams with an outgoing packet size of 64 samples. The receiver buffer holds two packets in the queue before decoding, meaning that a delay of three packets is encountered before the result can be heard. With a sampling rate of 44.1kHz, this translates to a packetization and receiving latency of  $3 \times (64/44.1) = 4.35\text{ms}$ . The network delay can be as low as 2ms if users are standing close to each other, and typically does not exceed 10ms for standard wireless applications. The sum of these latencies is in the order of 7-15ms.

Practical performance will, of course, depend on the wireless network being used and the number of streams transmitted. We have seen that high-traffic networks can become unstable, resulting in high jitter. In such situations it becomes necessary to increase packet size to help maintain packet rates at an acceptable level.

## 4. PROTOTYPE APPLICATION

Our initial prototyping devices, Gumstix, were chosen to provide: 1) wireless networking for bidirectional high-quality, low-latency audio and data streams, 2) local audio processing, 3) on-board device hosting for navigation and other types of USB or Bluetooth sensors, 4) minimal size/weight, and 5) Linux support. A more detailed explanation of our hardware infrastructure can be found in another publication [?], in particular, the method of Bluetooth communication between Gumstix and sensors.

To develop on these devices, a cross-compilation toolchain was needed that could produce binaries for the ARM-based 400MHz Gumstix processors (Marvell's PXA270). The first library that we needed to build was a version of Pure Data (Pd), which is used extensively for audio processing and control signal management by our Audioscape engine. Particularly, we used *Pure Data anywhere* (PDA), a special fixed-point version of Pd for use with the processors typically found on mobile devices [16]. Several externals needed to be built for PDA, including a customized version of the Open Sound Control (OSC) objects, where multicast support was added, and the *nstream* object, mentioned in Section 3.3. The latter was also specially designed to support both regular Pd and PDA, using sample conversion for interoperability between an Apple laptop, PC and Gumstix units.

We also supplied each user with an HP iPAQ, loaded with a customized application that could graphically represent their location on a map. This program was authored with HP Mediascape software [2], which supports the playback of audio, video, and even Flash based on user position. The most useful aspect of this software was the fact that we could use Flash XML Sockets to receive GPS locations of other participants and update the display accordingly. Although we used the Compact Flash GPS receivers with the

iPAQs for sending GPS data, the interface between Mediascape software and the Flash program running within it only allowed for updates at 2Hz, corresponding to a latency of at least 500ms before position-based audio changes were heard. The use of the GPSstix receiver, directly attached to the Gumstix processor, is highly recommended to anyone attempting to reproduce this work.




Figure 2: Mobile Audioscape architecture. Solid lines indicate audio streaming while dotted lines show transmission of control signals.

The resulting architecture is illustrated in Figure 2. Input audio streams are sent as mono signals to an Audioscape server on a nearby laptop. The server also receives all control data via OSC from the iPAQ devices and stores location information for each user. A spatialized rendering is computed, and stereo audio signals are sent back to the users. For all streams, we send audio with a sampling rate of 44.1 kHz and 16-bit samples.

In terms of network topology, wireless ad-hoc connections are used, allowing users to venture far away from buildings with access points (provided that the laptop server is moved as well). Due to the number of streams being transmitted, audio is sent with 256 samples per packet, which ensures an acceptable packet rate and reduces jitter on the network. The result is a latency of  $3 \times (256/44.1) = 17.4\text{ms}$  for packetization and a minimal network delay of about 2ms. However, since audio is sent to a central server for processing before being heard, these delays are actually encountered twice, for a total latency of approximately 40ms. This is well within the acceptable limit for typical musical performance, and was not noticed by users of the system.

The artistic application we designed was one where users could navigate through an overlaid virtual audio scene. Various sound loops exist at fixed locations, allowing users to congregate and jam with accompanying material. Several virtual volumetric regions are also located in the environment, allowing some users to escape within a sonically isolated area of the scene. Furthermore, each of these enclosed regions serve as a resonator of sorts, providing musical audio processing (e.g., delay, harmonization or reverb) to signals

played within. As soon as a player enters such a space, their sound is modified, and a new musical experience is encountered. Figure 3 shows two such performers, who have chosen to jam in a harmonized echo chamber. They are equipped with Gumstix and iPAQs, with both inobtrusively in their pockets.



Figure 3: Two participants jamming in a virtual echo chamber, which has been arbitrarily placed on the balcony of a building at the Banff Centre.

## 5. DISCUSSION

Approaching mobile music applications from the perspective of virtual overlaid environments, allows novel paradigms of artistic practice to be realized. The virtualization of performer and audience movement allows for interaction with sound and audio processing in a spatial fashion that leads to new types of interfaces and thus, new musical experiences.

We have presented the challenges associated with supporting multiple participants in such a system, including the need for accurate sensing technologies and network architectures that can support low latency communication in a scalable fashion. The prototype application that we developed was well-received by those who experimented with it, but many improvements still need to be made. The coarseness of resolution available in consumer-grade GPS technology is such that an application must span a wide area for it to be of any value. This is problematic, since the range of a WiFi network is much smaller, mandating redirection of signals through additional access points or OLSR peers. If all signals must first travel to a server for processing, then distant nodes will suffer from very large latency.

One solution is to distribute the state of the virtual scene to all client machines, and perform rendering locally on the mobile devices. For the prototype application that we developed, this would cut latency in half since audio signals would only need to travel from one device to another, without the need to return from a central processing server. Furthermore, this strategy would allow users to be completely free in terms of mobility, rather than in within contact with the server for basic functionality. However, for scenes of any moderate complexity, this demands much more processing power and memory than is currently available in consumer devices, and of course, the number of users will still be lim-

ited by the available network bandwidth required for peer-to-peer streaming.

A full investigation into distributing audio streams, state and computational load will be presented in future work, but for the moment we have provided a first step into the exploration of large-scale mobile audio environments. The multi-user nature of the system coupled with high-fidelity audio distribution provides a new domain for musical practice. We have already designed outdoor spaces for sonic investigation, and hope to perform and create novel musical interfaces in this new mobile context.

## 6. ACKNOWLEDGEMENTS

The authors wish to acknowledge the generous support of NSERC and Canada Council for the Arts, which have funded the research and artistic development described in this paper through their New Media Initiative. The prototype application described in Section 4 was produced in co-production with The Banff New Media Institute (BNMI) in Banff, Alberta, Canada. The authors would like to thank the participants of the locative media residency for facilitating the work and in particular, Duncan Speakman, who provided valuable assistance with the HP Mediascape software.

## 7. REFERENCES

- [1] Audioscape webpage. [www.audioscape.org](http://www.audioscape.org).
- [2] HP Mediascape website. [www.mscapers.com](http://www.mscapers.com).
- [3] Intersense. <http://www.intersense.com>.
- [4] The [murmer] project. [murmurtoronto.ca](http://murmurtoronto.ca).
- [5] Open computer vision library. <http://www.intel.com/technology/computing/opencv/>.
- [6] Pure Data. [www.puredata.info](http://www.puredata.info).
- [7] Webpage: 4Dmix3. [www.audioscape.org/twiki/bin/view/Audioscape/SAT4Dmix3](http://www.audioscape.org/twiki/bin/view/Audioscape/SAT4Dmix3).
- [8] R. Bajaj, S. L. Ranaweera, and D. P. Agrawal. GPS: Location-tracking technology. *Computer*, 35(4):92–94, 2002.
- [9] N. Bouillot. nJam user experiments: Enabling remote musical interaction from milliseconds to seconds. In *NIME '07: Proceedings of the 7th international conference on New interfaces for musical expression*, pages 142–147, New York, NY, USA, 2007. ACM.
- [10] N. Bouillot, M. Wozniowski, Z. Settel, and J. Cooperstock. Adaptive low latency audio streaming for mobile augmented reality applications. In *submitted to IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks (WOWMOM 2008)*, Newport Beach, USA, June 2008.
- [11] D. Brungart, B. Simpson, R. McKinley, A. Kordik, R. Dallman, and D. Ovenshire. The interaction between head-tracker latency, source duration, and response time in the localization of virtual sounds. In *Proceedings of the 2004 International Conference on Auditory Display*, 2004.
- [12] D. S. Brungart and A. J. Kordik. The detectability of headtracker latency in virtual audio displays. In *Proceedings of International conference on auditory displays (ICAD 2005)*, pages 37–42, 2005.
- [13] E. Chew, A. A. Sawchuk, R. Zimmerman, V. Stoyanova, I. Tosheff, C. Kyriakakis, C. Papadopoulos, A. R. J. Franois, and A. Volk. Distributed immersive performance. In *2004 Annual National Association of the Schools of Music (NASM 2004)*, San Diego, November 22 2004.
- [14] R. Etter. Implicit navigation with contextualized personal audio contents. In *Adjunct Proceedings of the Third International Conference on Pervasive Computing*, pages 43–49, 2005.
- [15] L. Gaye, R. Mazé, and L. E. Holmquist. Sonic city: the urban environment as a musical interface. In *NIME '03: Proceedings of the 2003 conference on New interfaces for musical expression*, pages 109–115, Singapore, Singapore, 2003. National University of Singapore.
- [16] G. Geiger. PDA: Real time signal processing and sound generation on handheld devices. In *Proceedings of International Computer Music Conference (ICMC)*, 2003.
- [17] Gumstix. [www.gumstix.com](http://www.gumstix.com).
- [18] P. Hipercom. RFC 3626, optimized link state routing protocol (OLSR), 2003.
- [19] K. K. Joseph Rozier and J. Donath. Hear & there: An augmented reality system of linked audio. In *Proceedings of International conference on auditory displays (ICAD)*, 2000.
- [20] S. Strachan, P. Eslambolchilar, R. Murray-Smith, S. Hughes, and S. O'Modhrain. GpsTunes: controlling navigation via audio feedback. In *MobileHCI '05: Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*, pages 275–278, New York, NY, USA, 2005. ACM.
- [21] A. Tanaka. Mobile music making. In *Proceedings of New Interfaces for Musical Interaction (NIME)*, 2004.
- [22] A. Tanaka and P. Gemeinboeck. A framework for spatial interaction in locative media. In *NIME '06: Proceedings of the 2006 conference on New interfaces for musical expression*, pages 26–30, Paris, France, 2006. IRCAM.
- [23] S. Wabnik, G. Schuller, J. Hirschfeld, and U. Krämer. Reduced bit rate ultra low delay audio coding. In *Proceedings of the 120th AES Convention*, May 2006.
- [24] M. Wing, A. Eklund, and L. Kellogg. Consumer-grade global positioning system (GPS) accuracy and reliability. *Journal of Forestry*, 103(4):169–173, 2005.
- [25] M. Wozniowski, Z. Settel, and J. R. Cooperstock. User-specific audio rendering and steerable sound for distributed virtual environments. In *Proceedings of International conference on auditory displays (ICAD)*, 2007.