

Syren - A Ship Based Location-Aware Audio Experience

Daniel Woo¹, Nick Mariette¹, Nigel Helyer² and Chris Rizos³

¹School of Computer Science and Engineering, The University of New South Wales, UNSW Sydney NSW 2052, Australia

²Sonic Objects, P.O. Box 195, Balmain NSW 2041, Australia.

³School of Surveying and Spatial Information Systems, The University of New South Wales, UNSW Sydney NSW 2052, Australia
e-mail: danielw@cse.unsw.edu.au Tel: + 61 2 9385 6495; Fax: +61 2 9385 5995

Received: 6 December 2004 / Accepted: 20 October 2005

Abstract. Syren, a location-based, multi-speaker augmented audio reality installation was presented as a shipboard exhibit at the 12th International Symposium on Electronic Art in August 2004. It was conceived as a continuous 3-day spatial audio experience that augments the landscape through the Baltic archipelago with location-based audio media, spatialised through a 12-channel speaker array. As the ship tracks between Helsinki, Mariehamn, Stockholm and Tallinn, listeners on the upper deck hear sounds that are perceived to originate from geographic features. Our custom GIS is derived from electronic nautical charting information that includes coastlines, buoys and beacons. A handheld GPS provides both position and direction data that was used by a software system to drive parameters of the spatial audio presentation. The sound production for the artwork was created using the custom application that enabled the artist to place sound media in relation to a real-world map. An important component to this software was the ability to audition the audio experience without ever taking the journey.

Keywords: Augmented audio reality, GPS applications, location-based media, Geographical Information Systems (GIS)

1 Introduction

Syren is an artwork that uses location-based spatial sound as the central user experience. It was presented onboard a passenger cruise ship on the Baltic Sea as part of the International Symposium on Electronic Art 2004. The journey was between Helsinki, Mariehamn, Stockholm and Tallinn. The nautical setting provided a unique opportunity to deploy a novel augmented audio reality

artwork, presenting listeners on the upper deck with a spatial sound experience. As the ship moved through the Baltic archipelago sounds placed by the artist would be presented to listeners through a multi-channel speaker array. The visual landscape was augmented with audio media that had been attached to physical landmarks. As the ship approached an island, the sounds attached to that island would become louder as if the island was the source of sound. As the island moved from bow to stern, the augmented sound moved and attenuated in concert with the visual stimulus of reality.

Syren is the first project to be produced by an art / science / engineering collaboration between sound artist Nigel Helyer, Daniel Woo (User Interface Design, Usability and Implementation) and Chris Rizos (GPS and Spatial Information Systems). The focus of the project is to better understand the user interface tools that will support the creation of artistically motivated location-based deployments. In this study, a ship-based setting is explored but in the future, the tools should work with a range of emerging mobile devices and wearable audio interfaces such as those described by Cohen (2002).

1.1 Prior Art

Not long after consumer handheld GPS receivers became available in the early 1990s, artists utilised GPS in a multitude of ways. Early GPS art explored the implications of the technology and its capability to record the position traces of movements in the world. Kurgan (1995) displayed points derived from a fixed GPS to illustrate the random nature of the GPS signal, and also produced images of the lines and the letters "MUSEU" recorded using a handheld GPS receiver. The Field-Works series of projects by Fujihata (1992) used raw GPS data to create digital images such as 'Impressing Velocity' in 1992-1994, and has explored the use of situated photographic images recorded with position

captured on a GPS device (eg. 'Tsumari' a workshop carried out in 2000). Many other artists have similarly used GPS in projects that annotate locations in the world with electronic media objects. This activity has encouraged the emergence of online artist communities concerned with situated media, such as the 'Locative Network'. Work by the group 'C5 Corp' is also notable with an aim to utilise "GIS and GPS technology for research into the epistemic function of human cooperation".

Another unique and pertinent example of GPS art is 'Sound Mappings' by Mott & Sosnin (1997), which is one of the earliest GPS art works using the medium of audio. In Sound Mappings, the urban landscape is used as the foundation of a real-time musical composition that responds to the GPS-tracked movements of participants.

Augmented reality (AR) superimposes synthetic representations of location-based information upon what we naturally perceive of the real world using our senses. Augmented audio reality specifically presents information via the sense of hearing using spatial audio. Positioning technology is a vital component of any AR system in order to locate the user in the world. Cohen (1994) identified GPS would be an appropriate positioning technology for audio augmented reality applications.

Location-based audio can be distinguished from audio AR in that it does not require content to be presented in a spatialised manner.

Another variant of audio AR is the project 'Hear and There' that used GPS as the outdoors positioning technology in a system that enabled users to annotate a courtyard space with sound samples (Rozier, Karahalios & Donath, 2000). The user can participate by generating or receiving audio content. A range of other audio AR projects and research exists that use positioning technologies other than GPS, such as radio frequency positioning and infrared.

Pedestrian user, location-based audio artworks have been implemented by Perry (2002), in 'Invisible Ideas', and Knowlton, Spellman & Hight, (2002) in '34 North, 118 West'. 'Invisible Ideas' is a collaborative site-specific, non-spatialised audio and visual piece set in the Boston Public Garden. It runs on a on a GPS-equipped Pocket PC handheld device. '34 North, 118 West' uses site-specific, non-spatialised audio presented in the streets of Los Angeles to tell a story about the local history of the railroad network.

GPS and spatial audio projects have appeared in Helyer's sound art pieces. Helyer and Rizos had worked together in 1999 when Helyer, the artist in residence at Lake Technology had developed Sonic Landscapes, a mobile virtual audio reality experience that was presented in St. Stephens graveyard in Newtown, Australia (Helyer

2003). As listeners moved through the graveyard they could hear spatialised sounds emerge from landmarks in a very realistic manner (Fig. 2). This included content relating to the occupants of graves. Dynamic content, such as a plane flying overhead and a ghost that followed the listener were included in the artwork. This was achieved using spatial audio rendering technologies developed by Lake Technology. The hardware consisted of a backpack mounted notebook computer, GPS with external antenna and digital compass head-mounted on stereo headphones (Fig. 1).



Fig. 1 Backpack worn by St Stephens visitor



Fig. 2 Visitor viewing a landmark in the graveyard

1.2 ISEA2004

The sea journey component proposed for the International Symposium on Electronic Art 2004 (ISEA2004) provided a novel arena to deploy a new location-based artwork. ISEA is a conference predominantly attended by practising artists, cultural and social theorists with a strong interest in electronic arts. Syren was designed for the helideck of the ship where listeners could walk within a region surrounded by an array of 12 speakers (Fig. 3). The speaker array provided the opportunity for many people to experience the spatial audio effect simultaneously without having to wear headphones.




Fig. 3 Plan of the upper forward deck alongside the helideck.

From the forward, upper deck vantage point, the audience could see the panoramic view of the ocean and hear an augmented audio environment in which sounds were perceived to originate from locations in the landscape (Fig. 4).



Fig. 4 Listeners would hear a sound to originate from distant landmarks

Being sea-based, the GPS-based artwork did not suffer from the “urban canyon” issues that commonly arise in city-based GPS applications due to trees, awnings and buildings. The open water setting provided a clear view of the available satellites. The most significant unknown was how well the digital compass module in the handheld GPS unit would operate when placed on board a large metal object such as a ship. Fortunately, the bearing could be calculated within the handheld GPS device based on the actual course of the ship when travelling above 10

knots. The digital compass proved to be inaccurate at slow speeds onboard the ship. The journey took place between Helsinki, Mariehamn, Stockholm, Mariehamn and Tallinn, taking approximately 41 hours, covering 520 nautical miles (963 kms). The artistic design challenge was to develop enough content to span the journey. Some aspects of the content were scripted but other aspects used collections of audio files that could be played randomly or in sequence. A granular synthesis effect (Roads, 2001) was used to create ambient background sounds that gave a rich but unusual presence, requiring only a small section of audio.

2 Technology

2.1 Hardware System

The hardware system (Fig. 5) consisted of a G5 2.0GHz dual processor computer running Mac OS X 10.3.4. A Garmin Etrex system was connected serially via a USB RS-232 interface to provide location data in standard NMEA format. A separate power supply was used to power the Garmin to avoid the need to change batteries over the 3-day journey.

Total analog audio output of up to 16 channels was provided using two external audio devices. A MOTU828 MkII audio output device provided 8 analog audio channels. An additional 8 analog channels was produced by optically connecting a Behringer ADA-8000 to the MOTU828. The analog channels from these devices were directly connected to 12 amplified speakers. Software was used to dynamically mix the audio over a number of speakers to achieve the spatial audio experience.

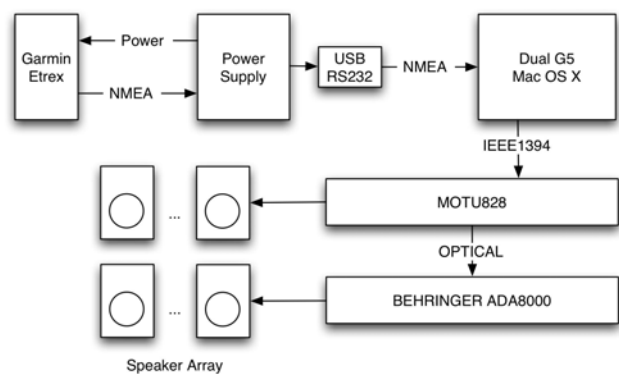


Fig. 5 Hardware Setup

2.2 Software System

The software system provided: 1) an accurate vector map representation of the Baltic Sea region; 2) an editing

system for the artist to place and audition sounds (and monitor the real time operation of the system); 3) a simple GIS, consisting of a database of audio regions, the sounds and audio effects associated with each region; 4) a GPS interface to provide both position and bearing information; 5) audio algorithms to determine loudness and angle of a sound source relative to the listener updated each second; and 6) audio rendering process to render the sound source so that it is mixed correctly for the speaker array.

2.2.1 Map representation

The map data representing coastlines and navigation features was derived from electronic nautical chart (ENC) data. ENC data was obtained from the Finnish, Swedish and Estonian maritime authorities and was made available in standard S57 format (International Hydrographic Bureau, 2000). This data is based on the ISO8211 file format which provides a method for accessing structured data. An open source library (Warmerdam 2004) was used to implement access to the ISO8211 file information. Two strategies were used to extract the data: 1) a simple extraction of the necessary features (eg. coastline outlines); and 2) a more complicated method based around ISO8211 data structures. The former turned out to be the most practical within the time constraints but reduced the total number of features we could immediately access from the data set.

2.2.2 Editing software

For editing purposes, a user interface tool, VectorMap (Fig. 6) was developed on Mac OS X using the Objective-C programming language and the Cocoa environment. VectorMap provided the artist with a digital map of the Baltic Sea with scales ranging from (1:5000 to 1:1500000) and a method for laying out circular regions containing one or more audio files. Each region and associated audio files were referred to as a sound stack.

A user-centred design approach was adopted to encourage a close working relationship with the sound artist. This was done to identify how the sound artist would use the tools rather than insisting that the artist learn how to use tools developed by engineers, for engineers. Many computer-based tools for audio design are based on the software developer's perspective, not an artist's. Our aim is to avoid this pitfall in design.

The artist needed to access a large audio library, audition the audio files and define regions in which one or more audio files could be heard. Listening is an important part of the sound artist's workflow and the ability to simulate the position and bearing of the ship a key feature of the

auditioning process. Previewing the audio content that was stored in the library was a requirement but there was also a need to understand the context of sounds presented in-situ. A map-based auditioning feature was mouse or keyboard driven allowing the ship position to be dragged with the mouse and the bearing modified with the keyboard. A variant on moving the ship was to use a metaphor of the ship being on the end of a string.

Circular regions were used over arbitrary polygon regions for reasons of simplicity. In physical terms it is straightforward to conceptualise a point source of audio radiating out over a circular region. Future versions of the software will support a range of tools for defining audio sound regions.

Each sound stack could be assigned a name, a display colour and the radius of the region could be adjusted. Variable opacity was used to clearly indicate to the designer the location of overlapping audio regions. When overlapping regions occurred, a context menu was provided to allow the selection of any specific region (by sound stack name) that contained the current mouse position.

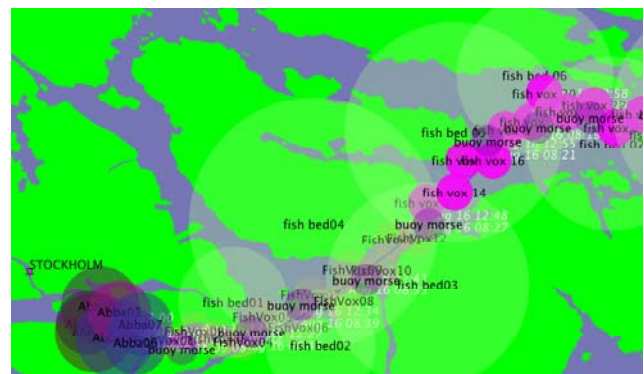


Fig. 6 Digital map user interface showing the audio design into Stockholm harbour.

Features that were requested but not implemented in the first version included support for audio layers, similar to image layers provided by graphical applications such as Adobe Photoshop, that could be enabled or disabled providing a mechanism for grouping audio regions with similar properties. The ability to search for regions by name was also suggested.

GPS track data from the shipping company was provided as a guide to indicate the path the ship was likely to take. The path was superimposed on the map indicating the expected times to reach the waypoints. Knowing the time of day that the ship was in a specific location was an important piece of information to assess whether people would be on deck or asleep. Design effort could be prioritised accordingly. Ideally, track data could be contained in a separate layer.

2.2.3 Database

The entire set of audio regions was defined in a custom database that essentially implemented a simple GIS appropriate for the project needs. The database was interrogated every second to compare whether the current position of the ship (the current position of the audition cursor or the actual position of the ship) was within one or more audio regions. The entire set of audio regions was contained in memory. No optimisation or caching was considered necessary for this prototype. This however provided rapid response when hit testing whether the region was to be processed or not.

The sound library audio content was not stored in this database, instead each audio region keeps a reference to the audio file locations on the hard drive.

2.2.4 GPS interface

The Garmin Etrex presented NMEA data serially at regular intervals of approximately one second. The device also provided compass information that was either calculated based on the current course (NMEA tag \$GPHDG) or directly determined via a magnetic digital compass module (Garmin NMEA extension tag \$HCHDG) embedded in the Garmin device. The Garmin simultaneously presented both pieces of direction data in the NMEA data stream. At slow speeds, the calculated heading by the device is inaccurate and the Garmin default threshold of 10 knots was used to determine whether to use the calculated or magnetic compass bearing.

The positioning aspect of the software system was designed so that it would be a simple task to switch between auditioning mode (where the user controls the boat position) and GPS mode (based on the NMEA data) without the underlying subsystems needing to know the source. This feature proved to be highly useful at slow speeds when the digital compass became inaccurate, and hence manual steering of the boat cursor was needed. It was useful for providing live demonstrations during the cruise, allowing manual positioning or rotating of the boat.

Code for GPS interfacing was based on the open source project Four Coordinates (Rosellini, 2003).

2.2.5 Audio lookup

In the VectorMap software, audio lookup is a process that occurs at regular intervals each time the GPS receiver generates a new position. The process manages a list of audible sounds that is based on their proximity to the current position. A search of the memory resident

database is carried out to identify the top twenty sound stacks that should be audible at any point in time. A region is audible if the position of the ship falls within the radius of the sound stack region. For each audible sound, the relative bearing and volume level is calculated and passed to the rendering process.

The audio lookup process is also responsible for managing the start and stop triggering of sounds when the ship has entered or exited a sound region.

2.2.6 Audio rendering

Audio signal processing was implemented using Pure Data (Pd), Miller Puckette's real-time, music and multimedia graphical programming environment (Puckette, 1996).. The rendering of the positional audio through the speaker array utilised vector-based amplitude panning (Pulkki, 1997). The number of simultaneous sound sources was arbitrarily set to 20.

Several audio effects were implemented to provide the artist with alternative ways to enhance the audio content. A single granular synthesis (Roads, 2001) processor was developed to generate spatialised, unrecognisable, random-sounding audio effects from given audio samples. This effect was used as ambient audio, and provided an easy method for creating passages of constantly varying, random sound within controlled timbral bounds. This provided a means to create large tracts of ambient sound content using relatively short sound files.

Sound source audio playback uses the Pd object 'readanysf~' by Black (2003). The audio system was controlled in real-time from the editing application via internal UDP packets containing data coded using the Open Sound Control protocol (Wright & Freed, 1997).

The original audio system included a 12.4 loudspeaker system that consisted of 12 loudspeakers arranged approximately in a circle, but were constrained by the layout of the deck. For editing purposes, a stereo audio output system was provided for headphone or desktop speaker use.

3 Discussion

3.1 Flexible Interfaces

Flexibility of the software tools was an important consideration for the project since there were many potential unknowns: this was the first time that such a project had been developed by this team in a geographic region that was unfamiliar to all of us. The artist was

concerned that we may overlook visible landmarks that whilst not obvious from the map data were glaringly obvious from upper deck of the boat. The system needed to support live updating of the audio database. This proved to be very useful when the ship approached Stockholm, docking at an area that was not expected from the preliminary data provided by the shipping company. The interface allowed rapid reconfiguration of the audio content to geographical area that had not been considered in the initial designs.

3.2 Artist On Board

Having an artist as an integral member of the team who was not fully aware of the underlying technical constraints of the system helped expose design constraints that would not have been found if this project was focussed only on engineering deliverables. As an example, we only ever anticipated 20-30 audio files being associated with one audio region. However, since we created an audio effect that could randomly play files from a selection of files it seemed quite plausible to place an entire collection of 400 audio files into a region and copy and paste multiple copies of the sound region where it was required. Whilst we did not hard code a maximum value, the additional time to access the data structure exposed subtle timing problems elsewhere in the application that indicated the need for threading.

As part of the VectorMap interface colour and opacity were used to help identify sound regions. In the final version of the piece, the artist had used colour in a highly artistic way producing graduations of colour to also show the progress along the journey. Throughout the exhibition, other artists commented positively on the visual aesthetics of the display. Hence, providing a simple feature like colour enabled another level of creativity and expression.

Artists, and in essence real users, push the boundaries of an application, which can produce a positive benefit to the overall outcomes of a project.

3.3 Tools not artworks

Another impression gained from art-centric conference attendees was the difference between tool design and desire to build an exhibit. In many cases, the art community follows the latter and purpose-built exhibits are created for display, literally as one-off productions. In the engineering community, tools and systems may be created but their applicability to artistic applications is often limited since the motivations for use are quite different. In Syren, the aim was to develop reusable tools that can be applied to a range of spatial audio and location-aware artworks and demonstrators.

3.4 Location-based Databases

Whilst the project has developed a custom database for storing the audio regions and references to audio media, the approach indicates what might be needed from future location-based services that are media focussed. From this project it is recognised that the designer may want to define arbitrary shaped regions for audio content not be limited to the basic shape of a circle. The database must support proximity queries allowing hit testing of a variety of regions and if real-time audio presentation is required, this process must respond quickly. A ship based project will not need to be as responsive as one in which the user is mobile and listening over headphones. Another issue for portable handheld devices is that of caching nearby regions and making predictions about what content might be required in the short term. Our current work has not focussed on this issue but it is identified as an important consideration for small, mobile memory devices that download media from a network system.

3.4 Future Directions

Our current course of investigation is to take what has been learned in deploying Syren to understand: 1) how the user interface design tools can be applied to a broader range of contexts such as campus guides and other forms of artworks; 2) the issues of platform scalability when considering handheld devices; and 3) how more robust database backend technologies can be applied to enable multi-user and collaborative applications.

Whilst the first version of the mapping interface is far from optimal, observation of the system in use by an artist has indicated plenty of opportunities for exploring how to solve navigation and searching issues that are related to the use of vector-based maps. Essentially, the editing system is a front-end interface to a GIS allowing the definition and interpretation of location-based information. In future work, this user interface will provide a test bed application to investigate navigation, search and usability issues that are relevant to all location-based information systems.

The map-based interface provides fundamental infrastructure to support a wide range of audio experiences that are location aware. Once the dependency on the physical speaker matrix is removed and a mobile solution is created, we can design and deliver audio to the device within any definable region where position can be sensed. The software tools have the potential to create pedestrian and vehicle audio experiences. The outcomes could include new forms of artworks, navigation, tourist guides and more generally, the contextual delivery of media depending on where people are currently located.

4 Conclusions

Syren is the first project milestone for the team and much of what we have learned from the Baltic deployment will be incorporated into the next iteration of our tool base. The ocean setting of the ISEA2004 conference provided a unique opportunity to undertake an ambitious, location-based spatial sound installation, which covered a journey of almost 1000km in around 41 hours. In comparison to other location-based artworks, Syren is vehicle- rather than pedestrian-based, covering a larger area than city-based pieces. The artwork is also an example of audio AR in that it uses spatialised audio to augment the environment. It did not enable audience input, as it did not support annotation of the piece with additional audio whilst in transit.

User-centred design helped focus our attention to the needs of the audio artist and assisted in the development of reusable tools, rather than a one-off exhibit. Artist involvement early in the development phase and throughout was a key factor to produce a highly productive editing environment. The approach that artists take when exploring tools can be highly creative and lateral in ways that engineers would not have readily foreseen, hence pushing the boundaries of possibility. Designers of augmented audio reality systems need to be aware of the needs of creative users who will be significant stakeholders in the future of location-based services.

Acknowledgements

This project is supported by an Australian Research Council Linkage Project in conjunction with the New Media Board of the Australia Council as part of the Synapse Initiative.

We wish to acknowledge the people and organisations whose contribution proved to be a valuable part of the Syren Project. ISEA Conference Staff and Technical Crew, Apple Computer Europe, Apple Developer Relations, Australia, Finnish Maritime Authority, Swedish Maritime Authority and Estonian Maritime Authority, Peter Mumford, Joel Barnes (SNAP Lab), Doug Kinlyside (Department of Lands), Natasha Leow, Bill Pepper (NIDA) and students from the Voice Studies course (2004) – Alistar, Emma, Geoffrey, Linda, Loretta and Tetsuya. Anthony Knittel and Ben Hudson.

References

- Azuma R. (1997): *A survey of augmented reality* In Presence: Teleoperators and Virtual Environments, Vol. 6, No. 4 (August 1997), pp. 355-385.
- Black A. (2003): *readanysf~* <http://aug.ment.org/readanysf/>, accessed May 2004
- Cohen M. (1994): *Augmented Audio Reality: Design for a Spatial Sound GPS PGS*. Center on Disabilities Virtual Reality Conference, 1994. California State University, Northridge <http://www.csun.edu/cod/conf/1994/proceeding/s/Aar~1.htm>, accessed November 13, 2004
- Cohen M. (2002): *A Survey of Emerging and Exotic Auditory Interfaces*, Proceedings of the 2002 International Conference on Auditory Display, Kyoto, Japan, 02-1 – 02-16.
- Fujihata M. (1992): *Field-Works projects website*, <http://www.field-works.net/>, accessed November 13, 2004
- Helyer N. (2003): *Sonic Landscapes*, <http://www.sonicobjects.com/sonic2/projectarchive/SONICLAND/SONICLAND.HTM>, accessed November 13, 2004
- International Hydrographic Bureau (2000): *IHO Transfer Standard For Digital Hydrographic Data*, Edition 3.1 – Monaco, November 2000
- Knowlton J.; Spellman N.; Hight J. (2002): *34 North, 118 West*, <http://34n118w.net/>, accessed November 13, 2004 from
- Kurgan L. (1995): *You Are Here: Museu* <http://www.princeton.edu/%7Ekurgan/urhere/html/intro.htm>, accessed November 13, 2004
- Mott I.; Sosnin J. (1997): *Sound Mapping: an assertion of place*. Proceedings of Interface 1997. <http://www.reverberant.com/SM/paper.htm>, accessed November 15, 2004
- Perry B. (2002): *The Invisible Ideas Project: Using Mobile Devices, Macromedia Flash, and Global Positioning Systems*. Macromedia Developer Center article, http://www.macromedia.com/devnet/devices/articles/invisible_ideas.html, accessed November 13, 2004
- Puckette M. (1996): *Pure Data*, Proceedings, International Computer Music Conference. San Francisco: International Computer Music Association, 269-272.
- Pulkki V. (1997): *Virtual sound source positioning using vector base amplitude panning*, Journal of the Audio Engineering Society, 45(6), 456-466.
- Roads C. (2001): *Microsound*. Cambridge, Massachusetts: MIT Press.
- Rosellini M. (2003): *Four Coordinates Framework*, <http://freshmeat.net/projects/fourcoordinates/>, accessed November 13, 2004
- Rozier J.; Karahalios K.; Donath J. (2000): *Hear & There: An Augmented Reality System of Linked Audio*. <http://citeseer.ist.psu.edu/rozier00hearthere.html>, accessed October 29, 2004
- Warmerdam F. (2004): *ISO8211Lib*, <http://gdal.velocet.ca/projects/iso8211/>, accessed November 13, 2004
- Wright M.; Freed A. (1997): *Open Sound Control: A New Protocol for Communicating with Sound Synthesizers*. ICMC 1997