Byzantium in Bing: Live Virtual Acoustics Employing Free Software

Fernando Lopez-Lezcano, Travis Skare, Michael J. Wilson, Jonathan S. Abel
CCRMA (Center for Computer Research in Music and Acoustics),
Stanford University
{nando|travissk|mwilson|abel}@ccrma.stanford.edu

Abstract

A Linux-based system for live auralization is described, and its use in recreating the reverberant acoustics of Hagia Sophia, Istanbul, for a Byzantine chant concert in the recently inaugurated Bing Concert Hall is detailed. The system employed 24 QSC full range loudspeakers and six subwoofers strategically placed about the hall, and used Countryman B2D hypercardioid microphones affixed to the singers' heads to provide dry, individual vocal signals. The vocals were processed by a custom-built Linux-based computer running Ardour2, jconvolver, jack, SuperCollider and Ambisonics among other free software to generate loudspeaker signals that, when imprinted with the acoustics of Bing, provided the wet portion of the Hagia Sophia simulation.

Keywords

impulse response, virtual acoustics, Ambisonics, auralization

1 Introduction

Acoustics are important to the experience of music: Singing in a large, stone cathedral evokes a much different response than the same singing in a small, wood frame recital hall. Acoustics are also important to the performance of music: Reverberation time can affect tempo, and room modes can influence pitch. Quite often, music written for a particular space works best when performed and experienced in that space.

As part of the "Icons of Sound" project, we are exploring the acoustics of Hagia Sophia, Istanbul, a nearly 1500-year-old World Heritage Site with marble floors and walls, 56-meter high dome and a reverberation time of over 10 seconds [1]. Hagia Sophia is presently a museum, and singing in the museum is not permitted. To better understand the aural experience of Hagia Sophia, we have attempted to synthesize the sound of Byzantine chant performed in Hagia Sophia.

Previous Icons of Sound acoustics and auralization work includes processing balloon pops recorded in Hagia Sophia into impulse responses of the space [2,3], and producing auralizations of Byzantine chant in a virtual Hagia Sophia [4]. The auralizations were accomplished by recording chant performed using headset microphones, so as to have separate dry tracks for each of the performer's vocals. While chanting, the microphone signals were processed using the estimated Hagia Sophia impulse responses, and played for the chanters over headphones to provide in real time a virtual sense of the performance space, while allowing dry vocal signals to be recorded. In post production, the recorded dry tracks were processed according to the estimated Hagia Sophia impulse responses to produce performance recordings in a simulated Hagia Sophia.

In this work, we describe a Linux-based system for live performance of Byzantine chant [10] in Stanford University's new Bing Concert Hall, modified to take on the acoustics of Hagia Sophia. Unlike virtual acoustic systems such as installed at McGill's CIRMMT [5], or LARES [6], any virtual acoustic system used at Bing must be able to be installed or removed within a few hours. For our Icons of Sound project, the system must also be able to handle the long reverberation times of Hagia Sophia.

The approach we take is to place two dozen loudspeakers in the hall, hanging many of them from the rigging points in the ceiling. Directional headset microphones are used to capture dry vocal signals, which are processed in a Linux machine to produce the needed acoustic enhancement.

This paper is organized as follows: Section 2 describes Listening Room and outdoor systems developed for synthesizing Hagia Sophia acoustics and listening at CCRMA. Sections 3 and 4 present the hardware and software components of the live virtual acoustics system used in Bing Concert Hall.

2 CCRMA-based systems

The auralization system used in the Bing Concert Hall performance was the result of a design and testing process lasting more than three years. The process included selecting and procuring the loudspeakers themselves, developing methods for diffusion and writing, testing and running the associated hardware and software. The real-time system was recently completed, and has been successfully used in concert several times.

The first public performance of one of the Icons of Sound auralizations was the Prokeimenon mix which was performed in the second night of the 2011 Transitions concerts [4]. This short piece was recorded in our small recital hall called the Stage by members of the Cappella Romana early music vocal group using the process briefly described above: headset microphones captured dry vocals which were convolved with Hagia Sophia impulse responses and fed back to the performers via headphones. In a post-production process, auralizations were created using an off-line convolution process to produce reverberated tracks for an Ardour session. The original mixing was done in our 3D Listening Room using its built in Ambisonics decoder and 22.4 loudspeaker configuration. The final mix was diffused outdoors using our (at the time) 16.4 system with height and a combination of 2D and 3D Ambisonics decoders.

In the second night of the 2012 edition of Transitions [7] we did the first test of a live performance in the virtual Hagia Sophia simulation for two chanters, John Kocolas and Konstantine Salmas. Two issues were being addressed, first, with critical help from Aaron Heller and Eric Benjamin [8], an expanded 24.4 diffusion system was implemented using the 3D outdoor diffusion system with proper Ambisonics decoding. Second, we found that Countryman B6 headset microphones could provide sufficient voice signal-to-reverberation ratios that the very wet acoustics of Hagia Sophia could be simulated without feedback. The concert was successful and proved the feasibility of using the system in a concert situation.

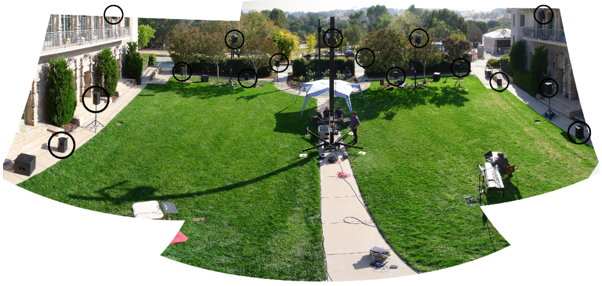

In December 2012 we had the opportunity to do a more complete test of the whole diffusion system in the framework of a two-day complete technical rehearsal in the Bing Concert Hall. It was the first time we had access to the Hall, and we installed a complete rigging of our 24.4 system in a full 3D dome.

We had not tested the system in an enclosed environment, and needed to verify that we could achieve the wet signal energy of Hagia Sophia without feedback. Using a group of seven Byzantine chanters from the choir of Holy Cross Church in Belmont, we were able to produce the needed levels of reverberation without feedback, provided head-mounted microphones were used.

3 Speakers and rigging

The final speaker array for the concert was comprised of 24 QSC HPR122i main speakers and 6 QSC HPR181 subwoofers. 10 main speakers were arranged around the upper terrace. The hall has a terraced design with the audience seated at different levels and surrounding the stage [18].

Because of the terraced design, regretfully that means that some of the audience would be necessarily close to a speaker and that is a real problem. To try to minimize that problem we kept an earlier arrangement of speakers mounted in equipment catwalks that are located between some of the "sails" of the concert hall. So we had, at “ear” level, two front speakers on stands, four speakers on each side of the Hall that were elevated (in the lower catwalks) and four more speakers in the sides and back on stands. This was a compromise between ideal location of speakers and trying to not get them to be too close to audience members.

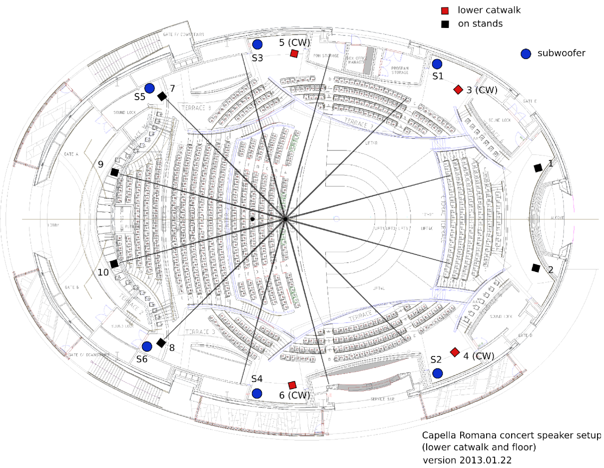

The speakers rigged from the ceiling were roughly arranged in two “rings”, a medium height one comprised of 10 speakers and a high ring of four speakers. The placement of the speakers was restricted by the spacing and availability of rigging points so it is not an exactly regular arrangement.

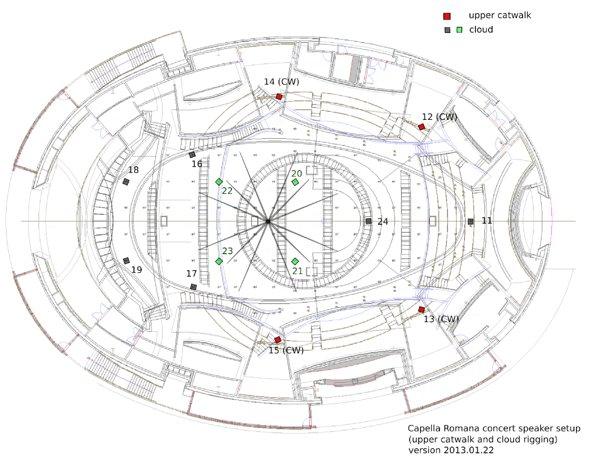

Aaron Heller and Eric Benjamin used their decoder design software to create an Ambisonics decoder tuned to the speaker arrangement we had in the Hall.

4 Signal routing

This section describes in some detail the signal routing of the whole diffusion setup.

4.1 Microphones

For the singers we used 16 Countryman B2D hypercardioid 2 mm diameter microphones with Sure wireless transmitters. Each singer had a microphone taped to his forehead. The wireless signals were picked up by antennas in the Hall which were connected to two racks of 8 wireless receivers each, located in the Amp/Patch Room to the side of the hall. From there they were patched into the hall's main Yamaha SL5 mixer which was used for level control and signal equalization. The 16 signals were then routed out of a dual ADAT link card into our sound diffusion workstation.

4.2 The workstation

Our workstation is a custom built no-fan workstation[11] currently equipped with a 6 core / 12 thread i7-3930K processor, 64G of RAM, an SSD system disk and a two disk mirrored RAID array for big audio and video files.

For this performance the computer had an RME RayDAT PCI Express audio interface which was slaved to the Yamaha digital mixer through Word Clock and received the microphone signals through two ADAT links coming from the Yamaha mixer. Jack 1.9.9 was used as the main audio connection and distribution system, and it was running during the performance at 256x2 and 48KHz. The preferred setting of 128x2 was actually not rock solid during tests so it was discarded, although it would have probably worked. The nature of the system being simulated did not require ultra low latency from the system (see below for the complex setup used).

4.3 Signal processing

All signal processing in the computer was done with Open Source and Free Software. An Ardour2 [13] session plus 4 instances of jconvolver [14] created the routing and main spatialization component of the Hagia Sophia virtual acoustics recreation running a total of 48 16-second-long convolutions.

Each of the 16 microphone signals has a pre-fader send to three jconvolver inputs. Inputs 1-4 are handled by the first jconvolver instance, 5-8 by the second, etc. 48 Ardour2 bus tracks each receive one mono jconvolver output. Each has a post-fader Ambisonics 3,3 insert to position the sound and output to the master bus. Virtual reverb sources were positioned manually via the Ambisonics plugin GUI inside Ardour. Additionally, a small amount of direct signal is mixed in via post-fader Ambisonics 3,3 plugins on the microphone input tracks. The project contains mix and edit groups for the dry (16) and wet (48) signals for ease of experimentation and mixing.

The resulting 16 channels of the 3rd order Ambisonics mix are then routed through Jack to the sound synthesis server of SuperCollider [16] (we were using the Supernova server running with 6 threads of parallelism, that corresponded to the 6 real cores of the machine).

SuperCollider received the Ambisonics signals and splits them into main speaker and subwoofer feeds using LR4 software crossover networks [9]. Those are sent out to two instances of Ambdec[15]. The first one decodes the main speaker feeds for our 24 QSC speaker dome, and uses coefficients calculated thanks to the help of Aaron Heller and Eric Benjamin decoder and optimization software [8]. The second decoder receives the subwoofer Ambisonics signals and decodes them with a standard hexagon decoder.

The output signals from the Ambdec decoders go back into SuperCollider, where a different set of SuperCollider instruments further process them to equalize all speaker and subwoofer feeds for delay and loudness (distances to all speakers and loudness are measured when the system is calibrated). At this point the signals for the 24 main speakers and 6 subwoofers are ready to be sent to the physical speakers.

4.4 Outputs

The workstation, which lives in the main mixing position of the Hall, is connected through a single ethernet cable to a Mamba AudioStreamer box that lives in the Amp Room. 30 of its 32 outputs are patched directly into the lines that are connected to each speaker or subwoofer.

The outputs of SuperCollider are connected to the inputs of jack-mamba[12], a small jack client that translates jack frames into UDP packets that the AudioStreamer box can understand.

To summarize: microphone signals are sent wirelessly to receivers that plug into the Yamaha digital mixer. Those signals are sent through ADAT links to the diffusion computer which routes them to Ardour2. Ardour mixes them, feeds four instances of jconvolver and receives 48 convolved signals from them which are spatialized using 3rd order Ambisonics panners. Ardour2 is connected to SuperCollider which does the crossovers for the subs. SuperCollider in turn feeds two Ambdec instances. Those are again sent through SuperCollider which compensates for latency and loudness and the final speaker feeds are fed into jack-mamba which sends them to the AudioStreamer box through ethernet UDP packets.

All of this running in realtime and providing a virtual Hagia Sophia auralization environment for the 15 singers of the Cappella Romana [17] vocal ensemble.

4.5 Control

A simple BCF2000 USB fader controller is connected to the workstation and three faders were programmed to control the 16 channel master Ardour bus (for overall level control), the DryGroup in ardour which controls the level of the dry spatialized signal being fed to the main mix and the WetGroup which directly controls the outputs of all 48 convolvers. Those three faders were used live during the performance to control the gain of dry and wet signals and overall loudness of the virtual hall. Fernando was controlling the mix during the performance.

The microphone signals were separately controlled by Doyuen Ko, a member of Wieslaw Woszczyk's team that was doing the recording of the performance.

5 Conclusions

The performance on February 1st 2013 demonstrated that the system can be used for real-time auralization of virtual environments, in this case the Hagia Sophia space. All the hardware components used are off-the shelf and relatively inexpensive, and all the software is all based on Open Source Free Software projects.

6 Acknowledgements

This project would not have been possible without the help and encouragement of many individuals. Chris Chafe (CCRMA's Director), Bissera Pentcheva (who worked with Jonathan Abel on this project for many years), Jenny Billfield (from Stanford Lively Arts), Aaron Heller for the timely delivery of optimized Ambisonics decoders, Sasha, Carr and Jay @ CCRMA and many others. The members of Cappella Romana and Byzantine chanters from the choir of Holy Cross Church in Belmont.

References

[1] Bissera V. Pentcheva, "Icons of Sound: Hagia Sophia and the Byzantine Choros," Chapter 2 in "The Sensual Icon: Space, Ritual, and the Senses in Byzantium," Penn State Press, 2010.

[2] Jonathan S. Abel, Nicholas J. Bryan, Patty P. Huang, Miriam Kolar, Bissera V. Pentcheva, "Estimating Room Impulse Responses from Recorded Balloon Pops," Convention Paper 8171, presented at the 129th Convention of the Audio Engineering Society, San Francisco, November 2010.

[3] Jonathan S. Abel, Nicholas J. Bryan, "Methods for Extending Room Impulse Responses Beyond Their Noise Floor," Convention Paper 8167, presented at the 129th Convention of the Audio Engineering Society, San Francisco, November 2010.

[4] Jonathan S. Abel, Bissera V. Pentcheva, Miriam R. Kolar, Mike J. Wilson, Nicholas J. Bryan, Patty P. Huang, Fernando Lopez-Lezcano and Cappella Romana, "Prokeimenon for the Feast of St. Basil (12th century)," from the concert "Transitions 2011, Night 1: Acousmatic Soundscapes under the stars," Center for Computer Research in Music and Acoustics, Stanford University, September 28, 2011.

[5] CIRMMT, Centre for Interdisciplinary Research in Music Media and Technology, http://www.cirmmt.mcgill.ca/

[6] David Griesinger, "Improving Room Acoustics Through Time-Variant Synthetic Reverberation," in Proc. AES 90th Convention, February 19–22, 1991.

[7] Jonathan S. Abel, Bissera V. Pentcheva, Miriam R. Kolar, Mike J. Wilson, Nicholas J. Bryan, Patty P. Huang, Fernando Lopez-Lezcano, John Kocolas and Konstantine Salmas, "Trisagion, Setting IX," from the concert "Transitions 2012, Night 2: Soundscapes under the stars," Center for Computer Research in Music and Acoustics, Stanford University, September, 2012.

[8] Aaron Heller, Eric Benjamin, Richard Lee, "A Toolkit for the Design of Ambisonic Decoders", Proceedings of LAC2012

[9] Linkwitz–Riley crossovers:

"http://en.wikipedia.org/wiki/Linkwitz%E2%80%93Riley_filter"

[10] Cappella Romana Concert:

"From Constantinople to California", Stanford Live, Bing Concert Hall.

http://live.stanford.edu/event.php?code=CAP1

[11] Fernando Lopez-Lezcano, "The Quest for Noiseless Computers", Proceedings of LAC2009

[12] Fernando Lopez-Lezcano, "From Jack to UDP packets to sound and back", Proceedings of LAC2012

[13] Ardour2

http://ardour.org

[14] Jconvolver

http://kokkinizita.linuxaudio.org/linuxaudio/

[15] Ambdec

http://kokkinizita.linuxaudio.org/linuxaudio/

[16] SuperCollider

http://supercollider.sourceforge.net/

[17] Cappella Romana,

"http://www.cappellaromana.org/"

[18] Bing Concert Hall at Stanford

"http://binghall.stanford.edu/about/"