ISEA2015
Proceedings of the 21st International Symposium on Electronic Art
ISSN: 2451-8611
ISBN: 978-1-910172-00-1

The Tomb of the Grammarian Lysias: Real-Time Performance and Crowd-Distributed Music Diffusion with Networked Mobile Devices

Ben Houge, Javier Sánchez

Berklee College of Music
Palau de les Arts Reina Sofia – Anexo Sur,
Av. Profesor Lopez Piñero, 1, 46013, Valencia, Spain
bhouge@berklee.edu, jsanchez5@berklee.edu

Abstract

"The Tomb of the Grammarian Lysias" is a setting of a poem by Constantine P. Cavafy for voice and audience members' mobile devices, composed by Ben Houge, based on software developed by Ben Houge and Javier Sánchez. During a performance of the work, a vocal soloist sings the poem in Greek, recording fragments of his or her voice using a custom application; these recordings are distributed wirelessly to the mobile devices of audience members for further processing and deployment, providing the crowd-distributed accompaniment to the soloist, with no other sound reinforcement required. The result is a uniquely portable and scalable performance environment in which the audience enables the work without directly interacting with it, representing an underexplored realm of app-based music performance.

This paper presents an overview of the work's genesis and antecedents, a description of the technology developed to enable the performance, and a discussion of its unique aspects and aesthetic ramifications. In closing we share some of the challenges related to presenting a piece that involves audience members' mobile devices, including a comparison of the work's two incarnations: as a native iOS app and as a web app using the Web Audio API.

Keywords

Mobile Apps, Networked Audio, Crowd-Sourced Performance, Granular Synthesis, Algorithmic Systems, Generative Music, Multichannel Sound, Greek Poetry, Microtonality, iOS, Web Audio API.

Genesis, History, and Previous Work

Constantine P. Cavafy (1863-1933) was an early Greek modernist poet who lived in Alexandria, writing in relative obscurity during his life, although his work has risen in prominence since his death, beginning with the advocacy of friends including E. M. Forster and champions like W. H. Auden [1]. Composer Ben Houge encountered Cavafy's poem "The Tomb of the Grammarian Lysias" on a trip to Greece in August 2013, and several aspects of the poem resonated with his work in real-time media and video games, prompting a desire to set it to music in a nonlinear treatment.

The poem is presented here in a new translation by Niko Paterakis, reproduced with kind permission:

> Very near to the right of the entrance to the library
> of Beirut, we buried wise Lysias,
> a grammarian. The spot is very well suited.
> We placed him near to those things that he may still
> remember there—commentaries, texts, technologies,
> variant readings, volumes filled with Hellenistic studies.
> And also this way, his tomb will be seen and honored
> by us, when we pass by the books.

The poem does not present a sequential, linear narrative; rather it describes a space, the final resting place of the fictional scholar Lysias. This kind of spatial organization suggests multiple, even infinite, possible trajectories through space, and this multivalence can be evoked in a musical treatment that incorporates statistical or stochastic processes at run time to create a dynamic sonic environment that varies from performance to performance. The poem goes on to list different types of texts found in a library, but not an ordered list, compounding a variable spatial organization with the notion of perusing a library's stacks of books (as one doesn't read a library from beginning to end). Finally, in several translations, the Greek term γραφές is rendered "variant readings," presenting a third enticement to set the poem in a way that maintains the idea of variation and extends it to live performance. These three observations as well as various personal associations provided the impetus to compose this music and to incorporate real-time, nonlinear processes.

The composer Ben Houge comes from a background in classical electroacoustic music via a nineteen-year career developing audio in the video game industry. The considerations brought to bear in this piece are informed by the exigencies of music composition for interactive digital media. Principal among these are twinned goals of allowing music to continue for an indeterminate duration (ideally using techniques more sophisticated and less fatiguing than a blunt loop) and allowing music to respond, transi-

tion, and evolve according to unpredictable events received as a result of real-time input. The medium of video games (and by extension video game music) is inherently indeterminate; once a system is made to allow user input, that system must be prepared to accommodate whatever agency has been afforded the user. When considered in these terms, the link to the spatially organized (rhizomic, and therefore variably navigable) themes of Cavafy's poem becomes apparent.

In addition, video games, especially 3D games, often deal with spatial organization in a more overt manner: sound sources positioned in 3D space may be encountered in infinite combinations as users variously navigate the virtual environment. Inspired by an understanding of spatially organized sound deployment in video games as a kind of virtual sound installation, Houge has developed several works for various configurations of multiple speakers, including the four-channel works *Lukou* (2009) and *Landscape with Water and Woodblocks* (2012), the six-channel composition *Kaleidoscope Music* (2009), the sixteen-channel incidental music for the dance performance *Please Be Seated* (2014), and the thirty-channel sound environment developed for a series of restaurant-based, audio-gustatory events called "food operas" (beginning in 2012), developed in collaboration with chef Jason Bond of Bondir restaurant in Cambridge, MA, USA, which presented an emergent, customized soundtrack for each diner during a five-course meal. [2]

So while this work does not contain an overt game element or virtual representation of space, the overlap with the concerns of video game music practice is considerable. In the discussion of the work's specific audio processing capabilities, additional connections to previous work will be presented.

The memorial nature of the poem as well as its language and cultural provenance suggested a link to the ritual music of Greek Orthodox Chant. This music is often sung not by a professionally trained choir, but by religious officiants, for whom the music's liturgical function takes precedence over performance practice considerations such as blend, precision, and intonation. The resulting sound possesses a rich heterophony that is well suited to a musical setting deployed across multiple mobile devices, transforming the audience into a kind of choir, spatially articulated, in which each member presents variations on common source material. This association with choral music practice is heightened by the fact that the sound being distributed is the multiplied voice of the soloist. [3]

The musical setting of this poem evokes a sense of solemnity and serenity, such as one might find at a funeral ceremony, again evoking the contours of traditional Greek Orthodox chant. Furthermore, in its application of indeterminate, real-time processes, this setting presents a series of static musical behaviors; while the piece has a finite (if statistically variable) length, the processes themselves could continue indefinitely, giving the music a sense of timelessness, fitting for a memorial. The work, with its strong links to the spatial aspects of sound installation practice, suggests an environment in which listeners could come and go at will and remain for an indefinite amount of time.

The setting is in just intonation, linking Cavafy to another prominent Greek from Alexandria, second century mathematician and theorist Ptolemy, who first articulated a musical tuning system based on varying small number frequency ratios. It has been observed that voices and other instruments capable of continuous pitch will tend towards just intonation when performing tonal music as a homogenous ensemble, and so this use of non-tempered tunings evokes once again the genres of choral music and Greek Orthodox chant. [4]

A final consideration in deciding to compose this setting for a soloist's voice deployed via audience members' mobile devices is the observation that recording, transmitting, and reproducing the human voice is the primary purpose for which mobile phones were designed, rendering this idiosyncratic medium uniquely suited for this type of setting.

Of course this is not the first work to exploit audience mobile phones.

Golan Levin, Gregory Shakar, and Scott Gibbons collaborated on *Dialtones (A Telesymphony)*, premiered in September 2001 at Ars Electronica in Linz, Austria, a famous early example that invited participants to register in advance to receive a pre-composed ring tone and sit in a designated seat in the concert hall; the piece was performed by dialing the phone numbers of the audience members according to a pre-composed scheme, resulting in a structured performance of about half an hour. Two hundred phones were registered, and up to sixty could be dialed at once. [5]

Jason Freeman's work investigates related issues, and his *Telephone Étude 1: Shakespeare Cuisinart* (2001) involved participants reciting text into a phone, where it was recorded, manipulated, and played back, although as a private, one-on-one experience. [6]

In some aspects, the format of our work is related to the Stanford Mobile Phone Orchestra (MoPhO), founded in 2008, which has spawned sibling ensembles in Helsinki, Michigan, and Hong Kong. While these ensembles focus on phones as instruments and coordinated performance, there is less emphasis on audience participation, resulting in fewer discrete sound sources, and the native acoustic properties of the phones are circumvented by the use of external speaker peripherals. [7] These Mobile Phone Orchestras can be considered portable descendants of the Princeton Laptop Orchestra (PLOrk) as described by Dan Trueman. [8]

The Collaborative Situated Media group at IRCAM has been working along similar lines, with their *Sound Checks* suite of Web Audio studies and the *Drop* composition they presented at the First Web Audio Conference in January 2015 (where the Web Audio version of our work was also premiered). [9]

Perhaps the researcher whose work mostly closely parallels to our own is Lonce Wyse at the National University

of Singapore, whose descriptions of the infrastructure, aesthetics, and challenges of his ADiffusion system mirror ours in many aspects. He also presented his work at the First Web Audio Conference and shared plans for a large-scale audience participation work in Singapore in the months ahead. [10]

Levin helpfully provides "An Informal Catalogue of Mobile Phone Performances, Installations and Artworks" on his website, comprised mostly of works from around the time of the premiere of *Dialtones*: http://www.flong.com/texts/lists/mobile_phone/. [11]

In January 2014, Houge established the Berklee Valencia App Choir, the first new ensemble in Berklee's new Music Technology Innovation program. This ensemble uses mobile devices to teach coding concepts with a tripartite focus: networked music systems (including audience devices), live sample manipulation, and on-the-fly score generation. The present work corresponds to this research trajectory.

Development

Javier Sánchez comes from a wide-ranging background in software development and computer-aided design. He brings extensive experience in app development, dating back to shortly after Apple opened up the iOS ecosystem to third party developers in 2009. He has linked his work in app development with electroacoustic research as a visiting scholar at Stanford University's Center for Computer Research in Music and Acoustics.

To develop the project, Houge began by composing the music for the vocalist and developing a prototype of the system in Max/MSP. Then he enumerated his goals and needs for the piece, and Sánchez proposed and outlined the architecture of the system framework, which Houge later expanded and filled in, programming in C, Objective C, and Python, with frequent input from Sánchez.

"The Tomb of the Grammarian Lysias" was premiered at the opening concert of the joint International Computer Music Conference/Sound and Music Computing Conference in Athens, Greece, on September 14, 2014, in the Ceremonial Hall of the University of Athens. The Chinese premiere of the piece took place on October 24, 2014, at Beijing's Central Conservatory of Music, as part of the Music Acoustica electronic music festival.

Houge later recast the work as a web application using HTML5, JavaScript, and the Web Audio API. This version was premiered at the First Web Audio Conference, co-hosted by IRCAM and Mozilla, January 26-28, 2015.

Performance Description

A rendering of the composition and score are available via the following links:

http://benhouge.com/resources/HougeTheTombOfTheGrammarianLysias.mp3

http://benhouge.com/resources/HougeTheTombOfTheGrammarianLysiasScore.pdf

As indicated above, the work exists in two formats: one based on a native iOS application and one based on a web application using the Web Audio API. We shall first describe the iOS version and then address how this work was adapted for the web, following the chronology of our work.

In the iOS incarnation, are three technical components to a performance of this work: the audience members' client app (for iPad, iPhone, or iPod Touch), available from the Apple App Store; the performer's iPad app; and a server running on a local wireless network. The performer uses the iPad app to record fragments of singing and to send messages to the server. The server broadcasts the data to the audience members' apps, which perform some local audio manipulation and play the sound using their built-in speakers. No other amplification or sound reinforcement is required, resulting in a highly portable piece that may be easily presented in unconventional spaces.

Timings are loose, there is no steady pulse, and metrical synchronization is not the desideratum; the local processes introduce random timing offsets while also accepting the unpredictable lag of data transfer over the network as a natural and inherent aspect of the piece.

A performance of the piece lasts a little less than five minutes.

Unlike the client app, the performer's app is not publicly available via the App Store. This app serves two primary functions: it can record fragments of the performer's voice during the performance and send them to the server, and it can send performance control messages to the server. Each audio file is associated with an index, corresponding to an allocated slot in each client app's memory, where it can be saved for subsequent deployment. There are eighteen available slots that can be filled with sample data from the performer's app. Audio is compressed using the AAC format before sending. Control messages trigger various audio deployment behaviors on the app, as described below. There is also a button to connect to the network and a button to stop all current processes. Several buttons serve more than one function (e.g., record and send a new fragment of audio while also triggering one or more audio behaviors on the client), in effect encoding the score of the piece into the app.

The client app presents a very simple interface, spread out over three different screens: the text of the poem, instructions, credits, and a connect button. The connect button instructs the application to search for a specific IP address on the network (currently hardcoded to the IP address 192.168.0.110, which on the router is a reserved DNS for the computer running the server, to streamline the setup process) and attempt to connect.

Once connected, the app can respond to network messages instructing it to download audio files from the server or to perform simple audio processes on files it has already downloaded. It has memory allocated to store eighteen

different audio files. There are two basic behaviors it can perform on this audio data. It can play the whole file back intermittently (choosing a number of repetitions, bounded by a minimum and maximum range, and a duration of time to pause between repetitions, also bounded by minimum and maximum values, as specified via incoming network messages). It can also perform a simple granular synthesis algorithm using the downloaded audio data as source material.

There are three variations of the app's granular synthesis behavior: a simple sustained tone, a string of several tones (specified by a set of possible frequency offsets, a minimum and maximum value bounding the number of tones to play, and a minimum and maximum value bounding the duration of each tone), and a variation on the string of tones in which pitch offsets are sorted such that they only move downward in pitch (resulting in a behavior consisting of the following steps: choose a tone from the available set at random; choose a tone between the current tone and the lowest possible tone; if the chosen tone is the lowest possible tone, and if a specified maximum time range has not elapsed, choose another pitch from the set at random and repeat the process). In all cases, the messages indicate an audio behavior of finite duration (even if the duration is variable and determined independently for each client); this is to preclude the possibility that any "note on" message might be sent without a corresponding "note off" message, resulting in stuck notes.

The iOS application launches upside down, encouraging users to intuitively invert their devices and point the speaker upwards for increased audibility.

There is no interface for audience member interaction, which sometimes surprises participants. Instead, the format of the piece can be thought of as a crowd-distributed speaker array. While the notion of audience participation in the performance of a composition via mobile applications is full of possibility, in this case, there was no meaningful way for the audience to contribute to the piece's progression; on the contrary it was feared that some trivial interactivity would detract from the focus on the piece.

The client app is designed to detect the device on which it is running and scale its functionality accordingly. For older, less powerful devices, fewer simultaneous sounds will be played; if a new sound is requested to play, but the device is already playing its maximum number of sounds, it will randomly decide whether to play the new sound or to continue playing the previous sound, promoting a more variegated heterophony throughout the space. In its first release, the app supports iOS7 and newer.

Both the client and performer app were written in Objective C. The simple granular synthesis algorithm is written using the low level Core Audio framework. The intermittent sounds rely on the high level AVFoundation library, although the intention is to move everything to Core Audio in a future update for more unified control.

The server is written in Python using the Twisted framework. Twisted implements a reactor design pattern. The bulk of the server consists of simple message parsing. When a message comes from the performer app, the server iterates through all connected clients and forwards the instruction to either download new data from the server (performing some checks to make sure all of the audio data has been successfully transferred) or perform some audio behavior on the client app.

The Web Audio API version of the piece is simpler in every regard. No advanced set-up is required from either the performer or the audience members; no local wireless network needs to be installed, and rather than downloading an app and connecting to a dedicated network, the audience can simply visit a web page URL projected on a screen. The performer and audience iOS applications are replaced by two different web pages. The communication system involves the performer app writing to a JSON file on a server (by means of a very brief PHP script), while the audience app regularly polls, looking for changes in the JSON file. (This leaves much to be desired, as the caching systems of various browsers often refuse to get a new JSON file, even if it has changed. We plan to replace this system with a more sophisticated and robust server architecture very soon.)

The version presented at the First Web Audio Conference did not involve real-time recording of the performer's voice; instead, recordings were made in advance, and when audience members visited the relevant web site, all of the necessary audio files were loaded into the corresponding buffers for future manipulation. Subsequently, Houge has developed a version of the work that does record and transmit audio in real-time on the web, making use of Matt Diamond's Recorderjs library, although this version has yet to be tested on a large scale. All of the other functionality of the piece is the same as for the original iOS version.

Technical Observations

We can draw several lessons from the three public performances of the work so far.

The premiere in Athens did not go off without a hitch. A challenge of being on the opening program was that there was little opportunity to request audience members to download the app prior to the show. The composer was obliged to recruit performers at the conference registration desk. As the venue did not have a public wireless network, one of the conference organizers supplied a private address that participants could use to download the app. Just prior to the performance of the piece, the composer (also the performer) announced the name and password of the local network that was used to run the piece, intending to remind the twenty or so participants who had been recruited earlier what they should do next. This was a bad idea. The network router, a D-Link DIR-655, was quickly overwhelmed with requests, presumably from audience members who thought that they would be able to connect to this network and somehow participate in the performance, not realizing that an app was required and that this local network was not connected to the internet. The performance started smoothly, but the composer was shortly dropped from the

server due to high traffic and was no longer able to control the piece. With the approval of one of the organizers, he restarted the server, and on the second try, the piece went flawlessly, but in the brief interim a somewhat confusing announcement was made in regard to the after-concert reception, and unfortunately many of the audience had already begun to leave.

The subsequent performance in Beijing, by contrast, went flawlessly from the beginning. For this performance, about twenty Central Conservatory students were recruited to download the app in advance, with a brief rehearsal prior to the concert. They were distributed through out the audience, and at the appropriate time in the program, they launched the app and connected to the network, enabling an ideal performance of the piece; these twenty performers represented about one fifth of the audience, resulting in a pleasing density of sound in the hall.

The version of the work premiered at the First Web Audio Conference, however, approached 100% audience participation. All that was required was for audience members to visit a URL that was displayed on a large screen at the front of the venue. An iOS device was no longer required, and by obviating the need for a local server, many more participants could connect. As a perhaps unnecessary precaution, the performer logged into a different wireless network from the rest of the audience (kindly lent by WAC keynote speaker and Mozilla software architect Paul Adenot).

Our observation from these performances as well as from our various tests and rehearsals is that, when the audience is involved, the set-up cannot be too simple. Even asking the audience to download an app and connect to a wireless network can be considered a complicated and error-prone request, especially compared to the ease of a web application. Also, maintaining a local network to run an on-site server represents an additional failure point that can be avoided with a web application.

Aesthetic Observations

The macro level structure of the piece is fixed, and each performance lasts about the same duration (with not much more variation than an a cappella singer might introduce from performance to performance). However, at the micro level, the piece is very open. Individual phrases play intermittently, with different numbers of repetitions and durations of pauses for each device, allowing them to interpenetrate, to use a favorite term of John Cage's. [12]

The processes are given a finite duration (or, in most cases, a range of durations) when they are set in motion, but the processes themselves are easily discernable as something that could continue indefinitely. As in much of the non-teleological music of Olivier Messiaen, there is a sense of a slice of eternity, of a music that could continue forever. In this way, the setting supports and reinforces the text, which describes a tomb, a memorial. Instead of a story with a determined beginning and ending, the poem presents a space in which the listener may linger indefinitely.

The result is a texture based on the singer's voice, varying in density, evocative of natural processes like rain or crickets chirping. These kinds of natural phenomena are often used to describe stochastic, granular textures, and in this case, the spatial aspect of these phenomena is also preserved, evoking, for example, the sound of the cicadas that Iannis Xenakis cites as a model (which were very active in Athens during the ICMC/SMC conference where the piece was premiered). [13] Diffusing this sound on multiple devices of similar playback capabilities throughout a performance space allows for a shift in perspective; the devices' built-in speakers are no longer merely poor speakers with limited frequency response and little dynamic range, an imperfect portal through which to imagine an ideal sound; instead these speakers become components in a larger spatial composition, instruments in their own right, and these attributes are embraced simply as the inherent qualities of a unique sound object, individually sited in space. In contrast to the common configuration of multiple speakers at regular intervals surrounding the audience, here the sound sources are interspersed among listeners. The performance takes on aspects of a sound installation.

The rich yet economically generated texture that results from multiple layers of intermittent sound has been explored by Ben Houge in several previous works, going back to the early 2000's, including the acoustic graphic score *A Reading from* _____ (2003), *Radiospace* (2004), and *Psalmus* (2005, also known as *Posuit flumina*). This type of organization also forms the basis of the music engine he developed for Ubisoft's real-time strategy video game *Tom Clancy's EndWar* (2008) and the thirty-channel interactive music system he designed for the food opera project mentioned earlier (beginning in 2012). [14]

While the small size of the speakers of audience members' mobile devices inevitably colors the sound significantly, it is important to note that these devices are designed primarily to convey human voices and are optimized for this type of signal, rendering a composition for voice and mobile devices particularly appropriate.

This configuration questions some of the most fundamental tenants of electro-acoustic performance. Since no sound reinforcement system is required, the system is highly portable, allowing for impromptu performances in unexpected locations. It is also scalable with the size and position of the audience. The stochastic organization of the sound materials is extended to include the playback apparatus itself.

A unique kind of community emerges from this arrangement. Through the act of activating and exploring the app on their personal devices, audience members become invested in the work and even responsible for successful performance. In many cases, a mobile device is a very intimate accessory, held in the hands, in close contact with the mouth or ear, so the degree of investment is greater than if users were simply handed a device for temporary use during a concert. Following the Paris performance, an audience member remarked that this was one of the rare events at which the audience wasn't distractedly fiddling

with their phones, sending messages or engaging with social media, since they were actively using their phones to enable the music.

However, we have rejected a persistent suggestion in relation to this work. Many people, when hearing the piece described, suggest inserting an element of user interactivity, providing the means to manipulate the sound as it is playing back. This is a valid approach, but one alien to the goals of this composition. To allow users to manipulate the sound would require either spending time to teach the piece to them in rehearsal or devising an interface that would constrain their behavior to an extent that it would fit within the objectives of the piece, rendering their actions trivial while distracting from the apprehension of the music.

Instead, this project identifies a little-explored middle ground between auditor in the traditional sense and performer in the digitally enabled sense; audience members operate in the field of the invested enabler.

Future Work

The portable, crowd-distributed, real-time sound deployment mechanism developed for Ben Houge's setting of "The Tomb of the Grammarian Lysias" is rich with suggestions for future work.

Most immediately, there is a desire to compose new work for this system, beginning with additional Cavafy settings (particularly other memorial poems or epitaphs, of which Cavafy composed several). In addition, the system could be expanded to include multiple performers recording their voices (or other sound sources) into their mobile devices, enriching the sonic resources available to the system. As of the time of this writing, Houge is preparing a setting of Elisa Gabbert's "Ornithological Blogpoem" for premiere at the first Vox Festival in Valencia, Spain, in June 2015, using a framework similar to the Web Audio API version of "The Tomb of the Grammarian Lysias."

One of the primary challenges of the work is to ensure that there are sufficient mobile devices connected to the system distributed throughout the audience to enable a satisfactory performance. It is expected that not every audience member will have a device, but due to the limited sound reproduction capabilities of typical mobile phones, a certain density, which varies with the size, configuration, and acoustical properties of the performance venue, is necessary to achieve if all audience members are to have a satisfactory listening experience. While this can be accomplished by planting participants in the audience, the crowd-distributed aspect of the work is only fully realized when audience members are able to participate freely with no advance notification.

By adapting our work as a web application using the Web Audio API, as opposed to a native application, we dramatically increase the number of audience members who are able to participate, without incurring the development overhead of support different native platforms. Audiences are no longer required to download a dedicated application in advance of the concert, but can simply be shown a web page URL at the time of the performance. The performer need not set up a dedicated network to run the piece; instead, the audience can connect to the venue's existing wireless network or use their personal data plans. This renders the work even more portable, allowing for performances in any space with no advance set-up.

Following the Athens premiere, one audience member commented that the iOS restriction (which was simply a function of development resources) could be considered elitist, denying participation to users of other devices. The Web Audio API version of the piece opens up participation to many more potential participants (although of course a mobile device of some kind is still required). However, a new set of challenges emerges, familiar to any web developer, which involves developing fallbacks to support as many different browsers and devices as possible. This affects not only Web Audio API support, but also audio format and compression settings. Other challenges mirror those enumerated by Lonce Wyse: audience members must actively keep their phones from going to sleep, app-based performances are dependent on the battery life of the phones, and a system that is dependent on a large number of simultaneous users is difficult to test. [10] Especially given the battery life dependency, works of long duration are challenging to execute. We observe rapid advancement in this area, and we anticipate being able to involve even more audience members as recent and emerging standards are more broadly adopted.

While the two systems developed for this composition exhibit many idiosyncratic qualities and pose some unique challenges, it should be expected that the mobile device ecosystem will continue to develop, resulting in increasing attach rate (already near ubiquity) for audience members, faster internet connections, broader adoption of new standards, and improved speaker quality. With this in mind, the authors believe that "The Tomb of the Grammarian Lysias" represents a promising paradigm for app-based, audience-enabled musical performance that is likely to increase in prominence as mobile technology continues to develop.

Acknowledgements

The authors thank Stephen Webber, program director for the Music Production, Technology, and Innovation department at Berklee College of Music's campus in Valencia, Spain, for his support of this project. Thanks to Anastasia Georgaki, Kostas Moschos, Cristos Carras, and everyone else involved in organizing the joint International Computer Music Conference/Sound and Music Computing Conference in Athens. Thanks to Benoit Granier and Zhang Xiaofu at the Central Conservatory for the invitation to perform in Beijing. Thanks to Victor Saiz and Samuel Goldszmidt at IRCAM for encouraging the development of the Web Audio API version of the piece. Thanks to Atlantis Books in Oia, on the Greek island of Santorini, for recommending the poems of Cavafy that provided the initial inspiration. And thanks especially to Niko Paterakis, who

provided extensive coaching on Greek language and pronunciation during the composition of this piece.

Bibliography

[1] Cavafy, C. P. *Selected Poems*. Translated by Avi Sharon. London: Penguin, 2008.

[2] Houge, Ben. "Food Opera: Merging Taste and Sound in Real Time." Paper presented at the Second Annual Workshop on Musical Metacreation, in conjunction with the Ninth Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, Northeastern University, Boston, Massachusetts, October 14-15, 2013.

[3] Choir of the Society for the Dissemination of National Music. *The Service of the Akathistos Hymn*. Simon Karas, director and soloist. Society for the Dissemination of National Music, SDNM CD 9, 1999, one compact disc.

[4] Partch, Harry. *Genesis of a Music*. New York: Da Capo Press, 1979.

[5] Traub, Peter. "Interview: Golan Levin." Accessed May 24, 2015. http://archive.turbulence.org/networked_music_review/2008/01/31/interview-golan-levin/

[6] Mirapaul, Matthew. "Coaxing On-Hold Melodies (Your Own) from a Phone." *New York Times*, June 25, 2001. http://www.nytimes.com/2001/06/25/arts/arts-online-coaxing-on-hold-melodies-your-own-from-a-phone.html

[7] The Stanford Mobile Phone Orchestra. "Stanford Mobile Phone Orchestra (MoPhO)." Accessed May 25, 2015. http://mopho.stanford.edu/

[8] Trueman, Dan. "Why a Laptop Orchestra?" *Organized Sound* 12/2 (August 2007): 171-179.

[9] Collaborative Situated Media. "CoSiMa at WAC '15." Accessed May 25, 2015. http://cosima.ircam.fr/2015/02/04/cosima-wac15/

[10] Wyse, Lonce. "Spatially Distributed Sound Computing and Rendering Using the Web Audio Platform." Paper presented at the First Web Audio Conference, Paris, France, January 26-28, 2015.

[11] Levin, Golan. "An Informal Catalogue of Mobile Phone Performances, Installations and Artworks." Accessed May 25, 2015. http://www.flong.com/texts/lists/mobile_phone/

[12] Cage, John. *Silence*. Hanover, NH: Wesleyan University Press, 1961.

[13] Xenakis, Iannis. *Formalized Music: Thought and Mathematics in Music*. Stuyvesant, NY: Pendragon Press, 1992.

[14] Houge, Ben. "Cell-based Music Organization in *Tom Clancy's EndWar*." Paper presented at the First Annual Workshop on Musical Metacreation, in conjunction with the Eighth Annual AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment, Stanford University, Palo Alto, California, October 9, 2012.

Author Biographies

Ben Houge has been developing audio for video games since 1996, including seven years at computer game pioneer Sierra in Seattle and four years at Ubisoft in Shanghai. Career highlights include his acclaimed string quartet score for *Arcanum: Of Steamworks & Magick Obscura* (2001) and the innovative cell-based music deployment system he designed for *Tom Clancy's EndWar* (2008). From 2004 to 2010, Ben lived in China, where he was active in the experimental sound scene. Previously he contributed Seattle's new music community as a founding member of the Stranger Genius Award-winning Seattle School collective. He has recently produced a series of "food operas," using video game techniques to score the indeterminate events of a five-course meal. Ben holds music degrees from St. Olaf College and the University of Washington and currently teaches in the Music Production, Technology, and Innovation department at Berklee College of Music's new campus in Valencia, Spain.

Javier Sánchez has been developing iOS apps since 2009. He is co-founder at Xculpture and Lingualia and has been involved in the development of iOS apps including Letsbonus, DressApp, Xculpture, iLoowi, Match My Music, and Multituch Me. From 2008 to 2011, Sánchez was a visiting scholar at the Center for Computer Research in Music and Acoustics at Stanford University, working at the Stanford Intermedia Performance Lab. He received his MS (1995) and PhD (2006) degrees in Mechanical Engineering from TECNUN, University of Navarra, Spain, where he focused on computational methods for the generation, manipulation and representation of tensile membrane structures in real time. His current work involves the relationship between art and technology and human computer interfaces, especially as related to industrial design, tensile membrane structures, and music. Sánchez current teaches app development in the Music Production, Technology, and Innovation department at Berklee College of Music's Valencia campus.